Optimal Tax Policy with Misreporting: Theory, and Evidence from Real Estate

Santosh Anagol Vimal Balasubramaniam

Benjamin B. Lockwood Tarun Ramadorai Antoine Uettwiller*

August 22, 2025

Abstract

We develop a model of optimal taxation and enforcement when policymakers seek both welfare maximization and "tax accuracy," where taxpayers remit the amount statutorily owed under truthful reporting. We characterize the optimal combination of tax rate and enforcement stringency. We apply the framework to the Mumbai real-estate market, a setting with widespread misreporting and a transparent enforcement instrument: government-specified guidance values set the minimum required tax base. Bunching evidence plus third party estimates of market values suggest under-reporting relative to market values, consistent with the government setting low guidance values to avoid the problem of overpayments. Structural estimates based on the observed level of guidance values suggest the government is willing to give up ₹2.12 in tax revenue to avoid charging taxpayers ₹1 more than their statutory liability.

*Anagol: Wharton School, University of Pennsylvania. Email: anagol@wharton.upenn.edu. Balasubramaniam: Queen Mary University of London and CEPR. Email: v.balasubramaniam@qmul.ac.uk. Lockwood: Wharton School, University of Pennsylvania and NBER. Email: ben.lockwood@wharton.upenn.edu. Ramadorai (Corresponding author): Imperial College London and CEPR. Email: t.ramadorai@imperial.ac.uk. Uettwiller: Queen Mary University of London. Email: a.uettwiller@qmul.ac.uk. This paper subsumes "A Bad Bunch: Asset Value Under-reporting in the Mumbai Real Estate Market". We thank Cristian Badarinza, Paul Carrillo, Edward Coulson, Sabyasachi Das, Raj Iyer, Venkatesh Panchapegesan, Tanner Regan, Shing-Yi Wang, Caroline Weber, and seminar participants at Imperial College London, UC San Diego, SITE Financial Regulation, Tilburg, Wharton, Johns-Hopkins SAIS, Princeton, WEFIDEV, 2022 Zurich Conference on Public Finance in Developing Countries, 2023 Syracuse-Chicago Webinar on Property Tax and Administration, 2024 MoFIR Conference on Banking, and the 2023 AREUEA National Conference for helpful comments. We thank Karthik Suresh, Karan Gulati, and Alexandru Zanca for able research assistance, and Raja Seetharaman at Propstack for data access. We acknowledge funding from Wharton Global Initiatives.

1 Introduction

The central goal of tax enforcement is to promote tax compliance and accuracy, so that taxpayers' actual remittances align with the amount that they legally owe. Inaccuracies arise in the form of "underpayments" if taxpayers pay less than their statutory obligations, e.g., due to evasion or avoidance activities. Inaccuracies may also take the form of "overpayments" in net taxes, e.g., if taxpayers fail to receive credits to which they are entitled, or if they are overbilled due to an error in tax administration. In this paper, we develop a model of optimal taxation and enforcement wherein tax authorities care both about maximizing welfare and about minimizing tax inaccuracies.

We apply the insights of this model to an empirical setting in a developing economy featuring property transaction taxes with self-reported transaction values. Many countries have attempted to control under-reporting by creating formulaic assessments of property value based on the physical location of properties, setting the tax base as the higher of this government-assessed value and the sales price reported by property buyers. Based on our review of transaction tax policies in the 82 largest cities in the world, 35 of these cities employ this specific system, and variants of this policy that involve some form of "model-based" appraisal of property values are even more widespread.[1] One such city is Mumbai, India, where we study the universe of residential property transactions between 2013 and 2022 using a large and granular administrative dataset. In this setting, home values serve as the basis for transaction taxes, capital gains taxes, and annual property taxes; overall, real estate transaction taxes constitute approximately 20% of state government revenues in India.[2]

The statutory requirement in the Mumbai setting that we study is that property purchasers report the true market transaction value to the government. If the buyer's reported value falls below the government assessed value (also known as the "guidance value"), then the tax base is set equal to the guidance value. An important mode of underpayment in this system is that buyers and sellers can collude to minimize transaction tax liabilities by under-reporting the true transaction price to the government.[3] The primary mode of overpayment in this system is that guidance values may

[1] A detailed spreadsheet of valuation systems for the top 82 cities of the world can be found here. We briefly discuss the broader context of cross-country variation in systems of property appraisals for tax purposes later in the paper.

[2] In Mumbai, there is a 5% stamp (transaction) tax, a 1% registration tax, and (small) property taxes are levied in certain sub-regions.

[3] The Indian tax administration (and anecdotal reports) discuss that the difference between the reported value and the true transaction price is often transferred from buyer to seller in currency notes, to avoid detection of tax evasion through the formal financial system. See Indian Department of Revenue's "White Paper on Black Money," 2012, `https://dor.gov.in/sites/default/files/FinalBlackMoney.pdf`.

be set higher than true market prices for some properties—a natural result of fundamental heterogeneity in land and property characteristics even in small geographic regions—forcing buyers to overpay in taxes, or to face costly and uncertain appeals processes.

How might governments optimally set tax enforcement policy, i.e. guidance values, in this environment to balance their twin objectives of maximizing welfare-weighted tax revenues and minimizing inaccuracy? Raising guidance values will increase revenues,[4] but will inevitably lead some individuals to overpay, thus raising inaccuracy. To understand and estimate this tradeoff, we develop an optimal enforcement model to estimate our key parameter of interest ψ—the amount of tax revenue the government is willing to forgo to avoid ₹1 of tax inaccuracy. To estimate this ψ parameter in the data, we take a revealed preference approach. We assume that the government has optimally set guidance values to jointly maximize tax revenues, redistributive concerns, and its preference for tax accuracy. To precisely quantify the government's preferences under this assumption, we develop and estimate a structural model of taxpayer reporting behavior. Using this model of reporting behavior, we estimate ψ as the preference for tax accuracy that maximizes the government's objective function and matches empirically observed guidance values.

Our structural model of agents' reporting incentives adapts and extends the classic tax evasion model developed by Allingham and Sandmo (1972) and Srinivasan (1973) to study how real estate buyers in India choose the value to report for transaction tax purposes. Buyers pay tax on the higher of their reported price and the guidance value, unless they file a successful appeal. The decision reflects a trade-off between saving on taxes and incurring the expected costs of audit, penalties for under-reporting, and appeals-related hassle factors. These expected costs depend on both the probability of being audited and the perceived likelihood that under-reporting will be detected. A key feature of our model is that audit risk changes discontinuously at the guidance value, generating bunching of reported transactions prices at that threshold. A buyer with a true market price less than the guidance value may also be forced to overpay, or to opt into a costly appeals procedure, so the model is well able to capture "crossovers" of both under- and over-reporting around that threshold.

Consistent with these modelled incentives, in our administrative data, we find prominent bunching of self-reported property transaction values at government-assessed guidance values. 7.3% of reported transaction values bunch within 1% of the guidance value, and an additional 13.9% report more than 1% below. Interestingly,

[4] This assumes small extensive margin elasticities, i.e., a small overall response of housing transactions to changes in transaction taxes; we confirm this is the case in our empirical work)

78.8% of transactions (corresponding to 26.5% of transaction tax revenues) are reported at 1% or more above the guidance value, suggesting that penalties and/or moral concerns impede a large part of the transacting population from simply reporting the minimum government value. We show similar bunching of reported transaction values at guidance values in data from Sao Paulo, Brazil, in Appendix Figure A1, suggesting that our methods are applicable elsewhere.[5]

While bunching itself is a likely smoking gun for tax inaccuracy (in the form of under-reporting), we are interested in the prevalence of both under- *and* overpayment of taxes, as our optimal tax framework suggests governments may have different preferences over these different types of inaccuracies. To estimate the prevalence of under- and overpayments we merge our administrative data on reported and guidance values with a third-party provided price dataset of new-builds developed and sold during our sample period.[6] These data serve as a proxy for market values, and are based on collecting pricing sheets and other marketing materials from developers. Unlike the distribution of reported property values around government-assessed guidance values, the distribution of this proxy for market values is smooth, with no bunching at guidance values. The visible difference between the distributions of self-reported and proxy market values constitutes the basis of our new technique to detect under- and overpayments of transaction taxes.

Using this new technique applied to the data over the sample period, we estimate under-reporting rates of 11%–30%; these average estimates come from underlying transactions that exhibit both under- and overreporting. When applied to Maharashtra state, this range roughly translates to annual losses for the state government of US$ 475mn to US$ 1.3bn lower tax revenues from property transactions in 2021—approximately 1% to 2.7% of total state government revenue. While substantial, these estimates are lower than the routine conjectures in Indian media and tax authority reports, which have for decades assumed that buyers part-pay in cash and substantially under-report real-estate valuations to launder undeclared income, such as cash earnings and bribes (so called "black money").[7] The desire to reduce such black money has motivated economically massive policy interventions, such as India's 2016 demonetization.

Relative to under-payment behavior, we find less evidence of over-payments. Even for the 14% of transactions that report a value less than the guidance value, we find

[5] We replicate this plot based on exhibits originally shown to us by Thiago Scot as part of their working paper Rocha, Scot and Feinmann (2023, mimeo).

[6] Multiple real estate analytics firms in India collect and sell these data, and we expect such data exists or will emerge in many developing country cities in the coming years.

[7] See, for example, Indian Department of Revenue's "White Paper on Black Money," 2012, https://dor.gov.in/sites/default/files/FinalBlackMoney.pdf.

most of these households are still under-paying taxes (even when their tax base is the guidance value) relative to our estimate of the market value. The fact that over-payments are rare is consistent with the government setting guidance values at relatively low levels, because it has an explicit preference for avoiding inaccuracy in the form of over-charging taxes.

Although bunching of reported values at government guidance values is consistent with under-reporting, it could also be consistent with truthful reporting for at least two reasons. First, while it is unlikely, infrequently updated government-assessed property values could be extremely accurate and timely estimates of true underlying transactions prices. Second, buyers and sellers might perfectly anchor transactions at guidance values.[8] If true underlying market values were observable, we could easily distinguish between these alternative explanations, and this motivates our use of the third-party price data as a proxy for true market value. While these data are purchased and used by banks, developers, and investors in the real estate space, we nevertheless consider the possibility that they are an imperfect proxy. We therefore adapt our approach and conduct several additional checks to account for measurement error.

First, we consider classical measurement error, where listing prices are noisy but unbiased estimates of true market values, and note that aggregation helps smooth out this noise. We show that if reporting is truthful and guidance values reflect true prices, aggregated reported and proxy values should align—but systematic under-reporting causes aggregated reported values to fall below proxy values, as we observe in the data. Second, in a robustness check we address the issue of households reporting values greater than our estimates of market values (behavior we believe is suggestive of measurement error). This form of measurement error takes a smooth form so we attempt to account for it in a parametric way. Our main estimates and results are similar under this measurement error correction. Third, we employ a more restricted dataset (70% of the full sample) of "exact matches" across reported transactions and third-party price data and confirm the robustness of our estimates. Fourth, we study revisions of government-assessed values (which occur in three of the nine years in our sample across multiple neighbourhoods) to confirm whether we are indeed uncovering inaccuracy in tax payments. Consistent with this explanation, we observe large spikes in the volume and value of registered transactions in the days and months immediately before scheduled guidance value changes, consistent with gaming behavior by agents rushing to register transactions right before these changes; we also find evi-

[8] While we are unaware of direct evidence that market prices anchor on guidance values, Genesove and Mayer (2001), Andersen, Badarinza, Liu, Marx and Ramadorai (2022), and Badarinza, Ramadorai, Siljander and Tripathy (2024) show that property sellers anchor on original purchase prices, and Garmaise and Moskowitz (2004) suggest that assessed values may be informative about market prices.

dence of agents back-dating transactions to pre-change dates to exploit lower guidance values.

Having verified the robustness of our empirical findings, we turn back to the main question that motivates this study. Intuitively, given that we observe (1) substantial bunching in reported values at government-assessed guidance values, (2) the distribution of guidance values lying substantially below the proxy market value distribution over much of its support, and (3) low estimated over-payments, we expect to estimate relatively high government preferences for avoiding tax inaccuracy. This is because with a lower preference for avoiding inaccuracy and a relatively higher weight on collecting tax revenues, the government would choose to raise guidance values to substantially increase revenue from taxpayers that report transaction values that exactly bunch at the guidance value.

To estimate ψ, we first map our structural model of taxpayer reporting to the data. To do so, we set some parameters in the model (such as the statutory penalty for detected under-reporting) at their observed values, and for the sake of parsimony, use single-parameter distributions to model underlying cross-taxpayer heterogeneity in other parameters (such as the hassle factor associated with filing appeals, and buyers' subjective beliefs about the probability of being audited). We find that the model delivers a good fit to the data, i.e., we are well-able to match the empirically observed distribution of reported values, counts, and (as an untargeted moment) the average under-reporting rate to the corresponding model-implied values using four estimated structural parameters.

Our second step is to estimate the weight that the government places on its inaccuracy minimization objective. We do so using the estimated misreporting aversion distribution from the first step to predict reported values, and vary the unobserved weight on the government's inaccuracy penalty to match model-implied guidance values with observed guidance values in the data across different areas of Mumbai. We estimate that the government has a high aversion to inaccuracy, and is particularly averse to overcharging taxpayers—quantitatively, the government is willing to forego ₹2.12 in tax revenues for each ₹1 of over-payment by taxpayers. To our knowledge, this is the first estimate of a government's preference parameter for avoiding tax inaccuracy.

In a final step, we validate the model and these estimated parameters by generating out-of-sample predictions of government guidance values from the model. We use these model predictions to forecast observed revisions to government guidance values in the data. More specifically, in areas where the model predicts that observed guidance values are too low (high), we later observe upward (downward) revisions to guidance values out of sample in the data. This exercise confirms the model's ability to capture the government's decision rule.

The paper is organized as follows. The remainder of this section reviews related literature. Section 2 presents our structural model of taxpayer reporting behavior as well as the government's optimal tax policy problem. Section 3 describes the institutional background on how property is valued for tax purposes in our setting, and introduces and summarizes the data. Section 4 documents empirical results on the observed distributions of reported values, guidance values, and our proxy for market values. Section 5 structurally estimates the model and validates it out-of-sample, and Section 6 concludes.

1.1 Related Literature

Our conceptual framework builds most directly on the literature on optimal tax administration and enforcement. Keen and Slemrod (2017), building on the early work of Sandmo (1981) and Mayshar (1991), presents a conceptual framework in which a policymaker maximizes social welfare by adjusting both tax rates and tax enforcement parameters such as audit probabilities. This framework identifies the *enforcement elasticity of tax revenue* as a key sufficient statistic, which quantifies the change in revenues—accounting for behavioral responses of taxpayers—of adjustments in enforcement policy. We build on this work by accounting for policymakers' potential desire for compliance and accuracy in tax administration, as distinct from pure revenue considerations.

Our empirical work relates to the literature on transaction taxes, which has mainly focused on advanced economies (e.g., Best and Kleven 2018, Kopczuk and Munroe 2015, Dachis, Duranton and Turner 2012), and has typically not estimated the importance of asset value under-reporting or associated optimal policy, presuming that third party reporting by mortgage lenders, real-estate agents, and other market participants eliminates the ability of buyers and sellers to under-report transaction prices. For example, Kopczuk and Munroe (2015) finds no evidence of evasion regarding a mansion transfer tax in New Jersey, and while Slemrod, Weber and Shan (2017) finds evidence of house price manipulation to avoid a higher average transaction tax rate in Washington D.C., they note that major tax evasion or avoidance behavior is unlikely in their setting given the relatively small change in average transaction tax rates studied. Sood (2020) highlights high transaction taxes as one component of land market frictions in India (and likely other developing countries) that may hinder productivity overall.

A small but growing literature studies under-reporting behavior in real estate transactions in China, though these papers do not analyze optimal enforcement policy in a context with under-reporting, which is our main focus. Fan, Wang and Zhang (2022), Agarwal, Li, Qin, Wu and Yan (2020), and Agarwal, Kuang, Wang and Yang (2020) compare reported values to data on underlying transaction prices collected by Chinese real-estate brokerages (the brokerages maintain records of true transaction prices because true prices serve as the basis for brokerage commissions). In the Mumbai setting, similar to many cities in developing countries, such administrative data are not recorded on true underlying transaction prices. Our approach for measuring under-reporting can be applied more broadly to detect under-reporting whenever the government sees reported and guidance values and can source (as we do) measures of market value or listing prices from analytics companies or online listings portals—in line with the broader agenda of technology-enabled improvements to developing-country tax collection (Okunogbe and Santoro 2022, Dzansi, Jensen, Lagakos and Telli 2024).

A related literature studies ongoing property taxation, finding that assessed values for property taxes can systematically diverge from recent transaction prices, with important distributional consequences (Avenancio-León and Howard 2022, Regan 2023). In this context, we are the first to analyse under-reporting in a system of government-assessed values. A key difference in our context is the statutory obligation for homeowners to report the true market value; in typical advanced economy property tax contexts, homeowners are required to pay tax on the government's assessed value, even if the assessed value differs from the market value (Amornsiripanitch (2020) reviews this literature).

In terms of possible remedies, Pomeranz and Vila-Belda (2019) survey research with tax authorities, focusing on policy interventions aimed at increasing tax revenues. To our knowledge this work has not studied real estate under-reporting nor the optimal setting of enforcement policies, especially in contexts where agents can choose to report at or above government-assessed values, though Casaburi and Troiano (2016) study the political economy consequences of an Italian national reform that aimed to force property owners to register their land so as to enter the tax base (an extreme form of asset value under-reporting is hiding property ownership from the government).[9]

Our analysis of property value under-reporting also connects to early analyses of black money in Indian real estate, which studied the government's pre-emptive purchase provision, under which the central government tax authority was allowed to purchase any property at 15% above the reported value, creating strong incentives for accurate reporting.[10] In statute, the government was supposed to randomly se-

[9] Our paper's focus on guidance values relates to studies on alternative minimum taxes and the evasion of corporate taxes - see Best and Kleven (2018) and Bukovina et al. (2025).

[10] This system appears to have been proposed in the economics literature by Harberger (1965), although Taiwan had a similar, largely unsuccessful, system implemented around the same time period

lect property transactions to determine whether to exercise this right, though most sources suggest the sampling was not conducted randomly. National Institute of Public Finance and Policy (1995) estimate 44.8% under-reporting using a small sample of Mumbai transactions under this policy,[11] and the same study conducts a survey of real estate brokers and concludes from this evidence that approximately 60% of true transaction values were under-reported (for earlier small-sample estimates see, Tandon 1987, Gopalakrishnan and Das-Gupta 1986). In a survey article on Indian transaction taxes, Panchapagesan and Karthik (2017) notes that there have been no more recent aggregate estimates of black money in Indian real estate transactions.

Finally, our (relatively low) estimate of the property value under-reporting rate is potentially useful for developing country studies on how neighborhood change, transportation infrastructure, and zoning reforms affect real estate prices, as these studies often use guidance property values in lieu of frequently unavailable market price data (Anagol, Ferreira and Rexer 2021, Tsivanidis 2019, Gechter and Tsivanidis 2020, Harari and Wong 2018).

2 Model

We model taxpayers who optimally report values that may differ from true underlying asset values. This model of misreporting serves as the foundation for our analysis of optimal tax policy when agents have incentives to misreport.

Canonical optimal taxation models typically operate under the assumption that tax authorities are solely concerned with maximizing social welfare (Ramsey 1927, Mirrlees 1971, 1986) when faced with individual agents who rationally optimize their economic decisions. Our goal here is to study how governments optimize policy when they are motivated *both* by welfare-maximization and "inaccuracy minimization" imperatives. This latter channel is a type of fairness criterion: Governments may not wish to obtain tax over-payments from taxpayers who truthfully report asset values, or to allow misreporting agents to substantially under-pay taxes relative to what they would owe if they truthfully reported.

Chang (2012). See Posner and Weyl (2019) for other examples of such "self-assessment" based mechanisms. A challenge to these systems is that those in charge of implementing the policy may be bribed to avoid exercising the government's right over certain properties.

[11] National Institute of Public Finance and Policy (1995) does not directly report the sample size for this estimate, however Table 3.1 in that study counts 46 properties purchased in Mumbai under this program.

2.1 A Model of Taxpayer Reporting Behavior

We adapt and extend the classic tax evasion model developed by Allingham and Sandmo (1972), Srinivasan (1973) and Yitzhaki (1979) to model individual taxpayer behavior. A taxpayer purchases an asset—in our empirical application, real estate—for market price m and then chooses the value r they report to the government, which could differ from m. The property has an associated government-assessed guidance value, which we denote by c (for "circle rate"; the nomenclature used in the empirical context, which we describe in detail later), and τ is the transaction tax rate. The transaction tax liability is therefore $\tau \times max(r, c)$, i.e., the tax base for the asset is the maximum of the reported value and the guidance value unless the individual chooses to appeal. If the appeal is granted, an individual with $r < c$ is permitted to pay τr in tax.

In practice, the guidance value c is computed on a property-specific basis using a coarse set of observables such as property size, some property attributes, and geographic location (as we later detail, the government divides Mumbai into "subzones" which share a common guidance value). To simplify notation in what follows, we consider, without loss of generality, the problem of a representative taxpayer with a specific c, r, and m. When we model the consequences of the government varying c in the analysis that follows, variations in this policy parameter appropriately rescale all guidance values within a subzone.

The taxpayer's optimal reporting amount after making a transaction trades off reductions in transaction taxes against increases in the probability of a costly audit (which potentially incurs a penalty for under-reporting).

The taxpayer's reporting decision is modelled as follows:

$$r^* = \arg\min_r \left\{ \tau r + \pi \cdot \left(t \cdot (1 + \beta(c - r)^+) + \left(1 - \frac{r}{m}\right)^\rho n\tau(m - r)^+ \right) \right\} \tag{1}$$

The buyer pays a tax on the reported transaction value, τr, with certainty. The parameter π, interpreted literally, represents the probability that the buyer is audited. We can also interpret π more broadly as the weight a buyer places on the aversive considerations involved in misreporting, such as their subjective probability of being audited, their inherent morality, or their fear of stigma.

If audited, the buyer incurs two types of costs. The first cost is a transaction verification cost, t, which can be interpreted as the cost of undergoing an audit in the case of involuntary verification, or as the cost of filing an appeal to reduce owed tax. An appeal may be filed by the taxpayer, for example, if the market transaction value m really lies below the guidance value. The cost t incorporates lawyer and accountant fees as well as any psychological costs or hassle factors associated with being audited or filing an appeal. If the transaction cost is incurred on account of an appeal where the buyer reports $r < c$, the cost t increases proportionally by $\beta(c-r)^+$, where $(c-r)^+ = max(0, c-r)$. Put differently, the costs of appeal scale with the extent to which r is lower than c, capturing the increasing difficulty of persuading the tax authorities to accept large gaps in statutorily owed taxes.

The second cost is the expected penalty owed to the authorities if the audit detects under-reporting. The buyer believes that conditional on an audit, under-reporting detection occurs with probability $\left(1-\frac{r}{m}\right)^\rho$, in which case he/she incurs an additional penalty of $n \cdot \tau(m-r)^+$, where n is the penalty multiplier ($n = 4$ is set by the authorities in the Mumbai context) and $\tau \cdot (m-r)^+ = \tau \cdot \max(0, m-r)$ is the evaded tax balance (note that there is no penalty if $m < r$).

The parameterization of the belief about under-reporting detection $\left(1-\frac{r}{m}\right)^\rho$ implies that this probability ranges from 0 (under truthful reporting with $r = m$) to 1 (under full misreporting with $r = 0$). The parameter $\rho \in [0, +\infty)$ controls the buyer's beliefs about under-reporting detection between these extremes, as Figure 1b shows. A low value of ρ implies that buyers perceive the detection probability to be high unless the reported r is close to the truthful transaction value m. A higher value of ρ implies that the perceived detection probability is close to zero even for those who report r well below the truthful value of m. Heterogeneity in ρ thus captures individual heterogeneity in beliefs about under-reporting detection.

We assume that the audit probability when $r \geq c$, denoted $\pi_{r\geq c}$, is bounded above by and less than 1. The lower bound of $\pi_{r\geq c}$ is governed by the penalty multiplier. To see this, consider the extreme case of $\rho = 0$ and $t = 0$. For an audit to matter at all to the agent even in this case, $n\pi_{r\geq c} > 1$, meaning that the probability $\pi_{r\geq c}$ must be at least $\frac{1}{n}$. If this is not satisfied, the tax savings from under-reporting always outweigh the penalty from under-reporting if the agent is caught, leading to a degenerate case in which all agents fully under-report.

We treat a taxpayer decision to report $r < c$ as a decision to trigger verification (i.e., to undergo a "voluntary audit"), meaning that by assumption, $\pi_{r<c} = 1$. This is to capture the reality in our setting that a taxpayer can reduce their tax base below c only by successfully appealing their case. Figure 1a plots model-implied π for values of r over and under c, and shows that there is a strong discrete incentive for agents in the model to report $r = c$ rather than $r < c$ (this is consistent with the bunching observed in the data, discussed in the next section). This incentive arises because a buyer that reports $r < c$ triggers an appeal, meaning that they do not pay τc (as they would if they reported $r = c$), but rather τr plus the verification cost $t(1+\beta(c-r)^+$, plus the expected penalty $(1-\frac{r}{m})^{\rho} n\tau(m-r)^{+}$ (which may be positive if they report $r < m$).[12]

The taxpayer in the model can choose to report $r = c$, $r > c$, or $r < c$; their payoffs are governed by the values of π that apply in each case. The solution to the minimization problem in equation (1) is:

$$r^* = m \times \left[1 - \left(\frac{1-(\beta\frac{t}{\tau})^{+}}{\pi n(1+\rho)}\right)^{1/\rho}\right], \tag{2}$$

where $(\beta\frac{t}{\tau})^{+}$ applies only when $r < c$. The buyer works out the r^* using equation (2) assuming $\pi_{r<c} = 1$, and also $\pi = \pi_{r\geq c}$, and chooses between reporting $r^*_{r<c}$, $r^*_{r>c}$, and $r^* = c$. The buyer then picks the value of r^* that minimizes their expected payout.[13]

The buyer's optimal choice of r^* thus depends on the parameters ρ, t, π, and β (not counting those set by the government, such as τ and n). To illustrate this dependence, the blue line in Figure 1c plots optimal reported transaction values normalized by the guidance value (i.e., r^*/c) for different values of ρ, assuming that $m = 2c$. The audit probability π jumps discontinuously from $\pi_{r\geq c}$ to $\pi_{r<c} = 1$ at the point $r = c$; the solid line in the figure shows that when $t = 0$, for a range of values of ρ, optimal reported values bunch exactly at the guidance value, i.e., $r^*/c = 1$ on the y-axis. The dashed line in the figure shows that when the transaction cost t is assumed positive, there is a noticeable increase in the range of values of ρ for which $r^* = c$.

Diffuse Bunching: The frictionless model discussed above delivers sharp bunching at c. It is possible that bunching in the data is not perfectly aligned with the guidance value, a common feature of many observed bunching distributions (see, e.g., Anagol et al. (2024)). To capture such diffuse bunching in the model, we extend the setup such that a buyer i can believe that the probability jumps from $\pi_{r\geq c}$ to $\pi_{r<c}$ at a value c_i that is slightly above c, i.e., $c_i = c + \epsilon_i$ where ϵ_i is assumed positive (this can be micro-founded as heterogeneity in buyer beliefs about the size of deviations from the guidance value the tax authority will tolerate before ordering an audit). We then have:

[12] While we have unsuccessfully sought administrative data on appeals and their outcomes, anecdotally, our understanding is that one reason buyers report $r < c$ is that they intend to appeal a guidance value that is set higher than the property's market value. We also note that it is possible in the model for transactors to report $r < c$ even when their market value $m > c$; we do observe this case in the data. Finally, it is theoretically possible for buyers to report $r < c$ but then choose to to pay τc rather than appeal—for example if they want the "truth" reported on the sale deed even though the system forces them to pay more.

[13] Technically, when $\pi = 1$, it is possible that $r^* > c$ minimizes the expected payout. However, we rule this case out by assumption, as $\pi = 1$ only occurs when $r < c$.

$$\pi(r,c) = \begin{cases} \pi_{r<c_i} = 1 & \text{with } r < c_i = c + \epsilon_i \\ \pi_{r\geq c_i} < 1 & \text{with } r \geq c_i = c + \epsilon_i \end{cases} \tag{3}$$

If $m > c$, then in the presence of this added friction, the discontinuity in π occurs at $c + \epsilon_i > c$. If $m < c$, then buyers choose either to report the truthful transaction value $r = m < c$ or to report a higher-than-truthful transaction value $r = c_i$ in order to avoid incurring the audit cost t.

Although t does not appear in equation (2) when $r > c$, it is particularly salient for buyers with $m < c$ but close to c. In other words, when buyers have a high t, but their true m is lower than c, they may choose to report at c to avoid incurring the t costs. Figure 1c plots the r/c values on the y-axis for various values of ρ, with and without t. As discussed earlier, even a small $t = 0.01$ increases the region for under-reporting, thus playing a role in the magnitude of bunching that arises from the mass coming from both under and over c.

The tax implications of misreporting are in reality more complex than this simple economic framework. For example, mortgages can reduce under-reporting because lenders base loan amounts on the property's officially reported value, incentivizing buyers to report more truthfully in order to access greater mortgage financing. When such external financing is needed, buyers face higher costs for under-reporting, such as sourcing expensive informal loans or settling for smaller properties, increasing the appeal of accurate or even inflated reporting to maximize formal credit access. Our method captures such motivations in the ρ parameter, in the sense that the incentive to report $r > c$ can be interpreted as a desire for mortgage credit rather than, or in addition to, a desire to avoid under-reporting detection.

There may be additional incentives for buyers to under-report to reduce ongoing taxes on property/asset values.[14] There may also be opposing incentives—for asset buyers to "over-report" to increase their current cost basis and reduce future capital gains taxes. We note that buyers who focus on minimizing the current transactions tax burden have under-reporting incentives that are well-aligned with sellers who wish to report depressed sale values to minimize capital gains taxes. However, financially-constrained buyers and forward-looking buyers that care about the capital gains tax basis more than transaction taxes may have opposing incentives to sellers, thus reducing opportunities for collusion to report low asset values. We do not consider such capital gains tax incentives for buyers in our model, primarily because the horizons

[14] One appealing feature of the institutional setting for our empirical work is that annual property taxes are only a function of the government's assessed value, not the reported value, so concerns about annual property taxes do not affect household's reporting behavior at the time of transaction. See https://housing.com/news/guide-paying-property-tax-mumbai/ for details.

over which housing transactions occur are very long, depressing the relative importance of this channel in determining government guidance value policy.

2.2 Optimal Government Policy

The government chooses values for the tax rate τ and the guidance value c—which we view as our policy instrument of interest—to maximize its objective, taking as given the taxpayer's optimization behavior over the reported value r^* described by the model above.[15]

In principle, a welfare-maximizing government selects these policies τ and c, as well as the myriad other parameters governing tax and regulatory policies to maximize some global objective function subject to a resource constraint. Rather than model this full optimization problem in all its complexity, we restrict attention to components of the objective that are sensitive to the selection of the guidance value c, implicitly holding fixed other dimensions of policy and specifying only the way in which marginal policy adjustments affect various dimensions of the government's objective.

The selection of the optimal guidance value c can be viewed as a choice of optimal enforcement activity along the lines of Mayshar (1991) and Keen and Slemrod (2017).[16] To represent the effects of marginally varying the guidance value c, the first relevant component is the effect on transaction tax revenues. Formally, indexing individuals by i distributed with measure $\nu(i)$, transaction tax revenues can be written simply as

$$R(\tau, c) = \int \tau \left(r_i^* + \phi n(m_i - r_i^*)^+\right) d\nu(i), \tag{4}$$

where ϕ represents the proportion of under-reported transaction values that are detected and recovered in the form of penalties. In practice, our empirical setting corresponds to one in which audits are very rare, and even recovered penalties are eroded by government-side transaction costs of verification and collection, so we consider $\phi = 0$ to be a reasonable approximation which we adopt hereafter. As a result, the revenue effect of marginally changing c is

$$R_c'(\tau, c) = \tau \int \frac{dr_i^*}{dc} d\nu(i), \tag{5}$$

As we will see in the empirical patterns in the next section, a prominent mass of taxpayers choose a "corner solution" with $r_i^* = c$, implying that for them $\frac{dr_i^*}{dc} = 1$, and that

[15] We also take as given the penalty multiplier n, as well as size and number of subzones, which the government also decides on. The smaller the subzone, the more tailored the government-assessed value can be, but increasing the number of subzones also comes at a cost.

[16] Viewed in this way, the first-order condition for optimal policy has as its analog equation (8) from Keen and Slemrod (2017).

revenue increases one-for-one with the guidance value. Many other taxpayers appear to exhibit partial misreporting, falling in the "fuzzy bunching" area above c but (presumably) below their true transaction value m_i.[17] These taxpayers' reporting response to a marginal change in c is governed by the structural model described previously, and simulated in our structural model to follow.

Balanced against the revenue considerations in equation (5) are the utility impacts on taxpayers from a marginal reform to c. In many optimal taxation settings, these utility impacts take a simple form: they are the mechanical financial cost of the reform for taxpayers—weighted by the taxpayers' welfare weight, usually representing their marginal utility of consumption—computed while holding fixed taxpayer behavior. The logic for this calculation stems from the envelope theorem: marginal adjustments in taxpayer behavior (such as earnings adjustments) have no first-order effects on their own utility because, starting from behavior that already maximizes taxpayer utility, such adjustments have no first-order effect on their utility. In this setting, however, that logic is complicated in two ways. First, taxpayers who "fully misreport" with $r_i = c$ appear to be at a corner solution, in which case the envelope theorem does not apply. Second, even for taxpayers at an interior solution with $c < r_i < m_i$, there is debate about whether their disutility from evasion should be given normative weight in the policymaker's objective function when c is raised, as the envelope theorem assumes. More generally, one might wish to apply a different normative weight to revenue raised from a compliant taxpayer (e.g., via a tax increase) than from a noncompliant taxpayer (e.g., via evasion).

To handle these nuances in a flexible but transparent way, we adopt a particularly simple specification in which we treat all of the funds raised from taxpayers in response to a c-reform as a weighted money-metric cost to them. This specification is very natural for taxpayers who fully misreport, in the sense that raising c mechanically raises funds from them as their fully misreported tax base increases one-for-one. The specification can also be understood as a reduced-form way of encoding the normatively ambiguous utility costs to partial-downward-misreporters with $c < r_i < m_i$, as it assumes a (flexibly-weighted) cost to them proportional to the additional tax they pay as they adjust r_i in response.

More formally, we assume that the welfare cost of raising a marginal rupee from individual i due to a small reform to c is equal to g_i. In this case, the welfare cost of a

[17] There is also an extensive margin elasticity for the tax authority to consider, i.e., the total number of asset transactions may fall if c is raised sufficiently high. This channel is represented in the revenue equation above by the possibility that r_i^* falls discretely to zero in response to a marginal increase in c. In our empirical application, we find no evidence of such an extensive margin response beyond very short-run time-shifting responses, and so we do not model a positive extensive margin elasticity (see Appendix D).

reform to c (excluding the revenue considerations above) is equal to

$$\tau \int g_i \times \frac{dr_i^*}{dc} d\nu(i) \tag{6}$$

In its most general form, the weight g_i may depend on many features of i's situation, including their marginal utility of consumption (as is conventional in models of optimal taxation) as well as other considerations, such as whether the taxpayer is evading taxes. In this most general interpretation, these are general weights as described in Saez and Stantcheva (2016).[18] For application in our context, we assume that these weights take a more specific form:

$$g_i = g(m_i) + \psi_{over} \cdot \mathbb{1}\{r_i < m_i\} - \psi_{under} \mathbb{1}\{r_i > m_i\} \tag{7}$$

where the component $g(m_i)$ represents i's marginal utility of consumption, which is assumed to depend on one's wealth, as proxied by the value of the property. The terms ψ_{over} and ψ_{under} allow for the government to place a different value on funds collected from taxpayers who are either over-paying or under-paying relative to τm_i. These weights can be viewed as a flexible reduced-form representation of a broad range of reasons a policymaker might value tax accuracy, such as the political costs of over-collecting from taxpayers relative to their true transaction amount, or the dynamic incentive costs of allowing underpayments to persist without penalty.

Combining equations (5) and (6) produces the net effect on the goverment's objective of a marginal reform to c:

$$\begin{aligned} \frac{d\mathcal{W}}{dc} &= \tau \int \frac{dr_i^*}{dc} \left(1 - g(m_i)\right) d\nu(i) \\ &\quad - \tau \int \frac{dr_i^*}{dc} \left(\psi_{over} \cdot \mathbb{1}\{r_i < m_i\} - \psi_{under} \mathbb{1}\{r_i > m_i\}\right) d\nu(i) \end{aligned} \tag{8}$$

This framework enables us to study optimal tax policy in the presence of misreporting, with individual taxpayers optimizing their reporting behavior, and the government balancing its twin objectives of raising revenue (accounting for the conventional welfare costs of their incidence on taxpayers) and minimizing inaccuracy of tax payments. In some of the analyses to follow, we consider the special case in which the government's aversion to inaccuracy is symmetric for over- and under-payment, in which case we simply use the notation $\psi := \psi_{over} = \psi_{under}$. Our goal is to illustrate how one can quantify the government's revealed for tax accuracy through this param-

[18] As shown by Sher (2024), such generalized weights must implicitly entail a global objective if they are to provide consistent rankings.

eter. We now turn to describing the institutional context, and empirically applying this framework to a rich dataset tracking the real estate market of Mumbai, India.

3 Institutional Background and Data

3.1 Institutional Background

Systems of property valuation for taxation purposes vary around the world, and can be broadly classified into two types. In the first type, taxation authorities generate "decentralized" (i.e., property-specific) assessed values, using a number of different approaches. In some jurisdictions, assessors determine property valuations using a combination of site visits and comparable analysis. In others, assessors determine the valuation by inputting features of the property (e.g., residential/commercial, square footage etc.) into a hedonic model.[19]

Physical assessment of individual properties can produce greater accuracy in determining market value, which is helpful given substantial unobserved heterogeneity in property quality even within small regions. Moreover, in decentralized assessment systems the guidance value is the statutory tax base; meaning that the owner has no reporting obligations. However, such individual assessments can be costly to implement and/or inaccurate given the scarcity of qualified assessors, and they can also be subject to manipulation, for example when assessors are bribed to lower assessments (Khan, Khwaja and Olken 2016). An alternative method of decentralized valuation is the so-called "self-assessment" method proposed by Harberger (1965), which encourages truthful self-reporting by giving the state the right to purchase the property at the property owner's self-assessed value.[20]

The Indian tax authorities (and many other jurisdictions, including, among others, Brazil, Colombia, Mexico, Thailand, Indonesia, Philippines, and New Zealand, see Appendix Table **??**) utilize an alternative "centralized" system of property valuation for taxation purposes. In such systems, the authorities assign location-specific (usually per-square-foot) valuations as a lower bound tax base for all properties in the physical location, and periodically update these valuations as market prices evolve. The statutory requirement is that owners report the true market value of their property, with

[19] For example, the Danish system of tax assessor valuation at different points in time adopted both property-specific and model-based property assessments (Andersen, Badarinza, Liu, Marx and Ramadorai 2022). In the United States, depending on the location and property type, local governments often conduct comparable-sale based assessments to determine the tax base. See, for example, this description of New York state property taxes, accessed February 2023.

[20] Chang (2012) argues that even with these incentives property owners in Taiwan greatly under-reported property values, because the probability of the state actually exercising the right was too low.

the tax base set as the maximum of the lower bound tax base (henceforth "guidance value") and the owner-reported value. The owner faces a penalty, typically a multiple of the amount of tax avoided if they report a value lower than the true market value (note that the true market value can be different from either or both of the guidance and reported values). Owners also generally have some form of (costly) recourse available to prove that lower valuations than the guidance value can be justified. Such centralization is more cost-effective than property-specific assessments, and can lower the probability of captured assessors. However, guidance values (especially when infrequently updated) can be inaccurate measures of market value.[21]

Conversations with market participants, and reports from regulators in the Indian centralized tax assessment system suggest that under-reporting follows a typical pattern.[22] The buyer (usually an individual) and seller (either an individual or a real estate developer) of an apartment agree on a transaction price. If this price is higher than the guidance value, to avoid taxes, they also agree to under-report the transaction price on the registration document (a frequent choice is to report exactly the guidance value; interviews with market participants suggest auditing is virtually non-existent as long as the reported value equals or exceeds the guidance value).[23] To prevent detection of the under-reporting by paper/digital trail, the gap between the transaction price and the reported price is paid physically in cash; an alternative is that part of the true transaction price can be misinvoiced as a higher "service charge" (this is especially prevalent in individual-developer contracts). This ensures that the bank records associated with the transaction are in agreement with the reported value on the registration, which is important, since buyers and sellers are required to report their tax-identification num-

[21] In parts of the U.S., tax authorities apply formulaic growth rates to historical assessed values, which is a similar approach (e.g., California's Proposition 13, see http://Santa-Clara-property-assessments accessed February 2023.)

[22] See, for example, https://timesofindia.indiatimes.com/blogs/law-street/black-money-does-the-devil-lie-in-real-estate/. A similar discussion of under-reporting real estate transactions in rural France is recounted in Mayle (2000): "This is the two-price purchase, and a typical example would work as follows: Monsieur Rivarel, a businessman in Aix, wishes to sell an old country house that he inherited. He wants a million francs. As it is not his principal residence, he will be liable for tax on the proceeds of the sale, a thought that causes him great distress. He therefore decides that the official, recorded price—the prix déclaré—will be 600,000 francs, and he will grit his teeth and pay tax on that. His consolation is that the balance of 400,000 francs will be paid in cash, under the table. This, as he will point out, is an affaire intéressante not only for him, but for the buyer, because the official fees and charges will be based on the lower, declared price. Voila! Everyone is happy."

[23] While the law states that guidance values should be formulaic, following centrally assigned guidance values, it is possible that the tax authority manually enters a valuation at their discretion. While this practice was not mentioned in our interviews with market participants, Comptroller and Auditor General of India (2016) discusses a few large transactions where the guidance value formulas were not followed. As we have administrative data on guidance values, our method also detects under-reporting arising from such inspector discretion.

bers to enable cross-verification with transactors' bank accounts. Typical methods of obtaining large sums of cash include accrued currency from business operations, a sequence of smaller withdrawals from bank accounts over time, or writing a check as a "gift" to a friend or relative in exchange for the cash. In some cases, cash funds can also be sourced or earned completely outside the tax net; this is referred to as "black money" in the Indian context.

3.2 Data

While market participants state that under-reporting is common, the general sentiment is that market prices are well understood, especially by developers and real estate brokers. Brokers are even known to quote market prices based on the appetite for under-reporting (i.e., a lower overall price for a higher cash share). However, it is generally not an easy task (and one to our knowledge that has not been undertaken) for the tax authority to acquire systematic data on market prices (by interrogating brokers and/or developers, or scraping and accurately matching data from online housing listings) and conduct the kinds of comparisons that we undertake in this paper. We now describe the specific data sources that we compile and merge for this purpose.

3.2.1 Registrar Data

The first dataset that we employ comprises real estate transaction registration documents from the Inspector General of Registration and Controller of Stamps (IGR), Department of Revenue, Government of Maharashtra, India. For our analysis, the important information in these documents is: 1) the reported property value; 2) the guidance value; 3) the transaction tax paid; 4) the property's floor space area; 5) information about buyer and seller type (corporate or individual); 6) the transaction date; 7) the registration date. These data are publicly available from the registrar's website, and cover all available transactions between 2013 and 2018. See Appendix Section B.1 for a complete description of this data and example documents.

We augment the IGR administrative data with data provided by Propstack Analytics which cover the period from 2013 to 2022. Propstack (`https://www.propstack.com/`) is a for-profit real estate firm that uses transactions to provide pricing and ownership information via its Zapkey (`https://www.zapkey.com/`) data platform. In addition to the IGR variables described above, Propstack Analytics also provides us: 1) an indicator for whether a property was sold by the developer directly (a "primary" sale) or sold by an individual; 2) the number of buyers; 3) the unit number; 4) the floor number of the apartment; 5) the name of the real estate project associated with the transaction; and 6) the latitude and longitude of each project location in Mumbai. Propstack covers all transactions reported in the IGR admin data, thus resulting in no loss of information for our analysis. We further confirm that the overlapping data from the two sources are identical; a 10% random sample of Propstack transactions between 2013–2018 perfectly match IGR reports. Figure A2 presents a comparison of aggregate counts of transactions and tax revenue from the IGR data, Propstack Analytics, and the aggregate numbers of transactions reported by IGR for the Mumbai Metropolitan Region—a region larger than the coverage of our sample but the closest level of aggregation for comparison.[24]

Between IGR and Propstack Analytics, the data cover the universe of registered residential real estate transactions in Mumbai and Mumbai suburban areas from 2013–2022 worth US$106.92 billion. This region is the most important metropolitan area in the state of Maharashtra, a state that generates a quarter of India's GDP. Our region of study remits approximately 30% of the state's total stamp tax revenues.[25] There is also comprehensive spatial information available for this region, which we use in our empirical analysis.

3.2.2 Propequity Data

Our independent source of market price data (i.e., the proxy p measuring m) comes from Propequity (https://www.propequity.in/), a real estate analytics firm that maintains a subscription real estate information portal for the Indian real estate market. Propequity is a for-profit analytics firm that primarily earns revenue by selling access to its data products. The subscribers are real estate public and private equity investors, banks and real estate developers. The primary use case is to understand trends in local prices and quantities for new residential projects being developed.

Propequity aims to provide data on all new real estate projects in India with potential revenues over 10 million rupees (roughly US$ 200,000), with coverage varying over locations. Over the time period 2013 to 2022 this dataset includes information on approximately 11,930 real estate projects (each such project has multiple apartment buildings which in turn have multiple apartment units) from the Mumbai and Mumbai suburban regions. For each project, we observe a number of characteristics: longitude, latitude, a masked developer ID, the number and format of apartment units and

[24] The government guidance value is determined by multiplying the guidance value (set on a per square meter basis for a given sub-zone × year within the city) by the area of the property. Additional adjustments to the guidance value are made based on other features, such as the floor on which the property is located or whether a parking spot is included with the property.

[25] Region-wise stamp tax revenue sourced from https://igrmaharashtra.gov.in/dashboard_Data_ArticlewiseAndYearwise.aspx?GvData=maharashtra.

amenity information, date project units started being sold, date project construction was completed, luxury status, and an estimated current sales price of the apartments in the project (the main variable for our purposes), in addition to a few other features. The quarterly price data are reported as the price per square foot for a "base" level apartment in the building (i.e., excluding optional amenities like parking spaces, higher floor levels, etc.).

These price data are sourced through two major methods: 1) physical visits to developers' sales offices to collect pricing sheets for projects and 2) collecting developer-emailed advertisements of projects, which report prices.[26] Developers typically market their apartments at a per-apartment price. The data provider converts this to a per square foot price using the developer's reported "carpet" area per apartment, which is the area of usable space within the apartment including interior walls, but excluding exterior walls, outdoor spaces such as balconies, and any public spaces within the building. The provider uses these data to conduct valuations for banks that make mortgages, and note that in this context they are asked to provide an estimate of the total value of the apartment, including both the reported value and any under-reporting to avoid the transaction tax. [27]

These price data serve as our proxy p for m, the true market value of transactions in the administrative data. We match Propstack transaction level data to Propequity price data, using the name of the project in which the transaction occurs, and the location (by latitude and longitude) every quarter. Of the total of 260,614 transactions recorded in Mumbai and Mumbai suburban regions between 2013 and 2022, 60.01% or 156,645 Propstack transactions are exact project-level matches to Propequity. We match the remaining 40% of transactions to the most geographically proximate Propequity project. Figure A3 documents the match quality in the data. Overall, 95% of all registered transactions are matched to Propequity transactions within 500 meters of the latitude-longitude of the purchase transaction, and we eliminate from our analysis successful matches that are further than 1 kilometer away from the location of the associated Propstack transaction. Figure A4 shows the spatial distribution of the transactions in our final sample for study. While the overall match quality is high, we carefully investigate effects of measurement error in p in greater detail later in the paper.

Table 1 presents means, medians, and counts by year for our primary analysis variables. 71% of transactions are sales made by a corporate entity (typically a real estate

[26] An example pricing sheet is given A7.

[27] Propequity also reports an estimated number of units sold within the building in a given quarter. As we have administrative data on sales from the registration documents, we do not use this information in our main analysis.

developer), with the rest made by individuals. The average property is 76 square meters (818 square feet) in size. The average reported value r over the sample is US$ 321,770. The average government guidance value c is 19% lower, at US$ 261,860, but the average p value is higher, at US$ 367,290. Based on these raw averages a simple calculation delivers an under-reporting rate of 12.4%; we approach this more rigorously below.

Appendix Figure A5 shows a binned scatter plot of p, r and guidance values c from our main sample of 260,614 transactions. The p and r values are highly correlated with c, suggesting that c values are set to match geographic variation in market prices. Reported values r are lower on average than estimated market values p, but also strongly correlated with them. As mentioned earlier, the positive relationship between p and c lends credence to the idea that guidance values can be useful proxies for market values in developing country cities where high quality individual transaction market price data is not available (e.g., see Anagol, Ferreira and Rexer 2021, Tsivanidis 2019, Gechter and Tsivanidis 2020, Harari and Wong 2018).

4 Empirical Results

In this section, we document bunching and compute a measure of property value under-reporting using the data. Figure 2a plots the empirical distribution of $\frac{r-c}{c}$ around zero (the blue triangles). The figure shows counts of transactions within 2% bins, with the bin around zero ranging from -1% to +1%. All three figures in Figure 2 also include our structural model fits, but we postpone discussion of these to the next Section.

The plot reveals that 13.9% of households report more than 1% below the government-assessed value; the average (median) r for this group is 24.9% (14.7%) lower than c. On the face of it, this is evidence that c imperfectly tracks m, since a non-negligible fraction of households is willing to pay verification costs to certify $m < c$. Second, there is visible bunching of reported r at the government-assessed c.[28] Figure 2 panel (b) shows the aggregate tax-base (i.e. the aggregation of $\max[r, c]$) value of transactions within $\frac{r-c}{c}$ bins, reported in billions of Indian rupees. The blue triangles here confirm that a large fraction of the tax base comes from transactions that report/bunch exactly at the guidance value, further motivating our structural analysis of how guidance values are set.

The red triangles in Figure 2 panel (a) show how the proxy for market value p

[28] Appendix Figure A6 shows transaction counts with one-tenth of 1% bins $\frac{r-c}{c}$ bins, showing that the bunching of transactions at the guidance value is distinct and sharply identified.

is distributed around c. In contrast with the distribution of reported values, the $\frac{p-c}{c}$ distribution is centered at roughly 30%, and appears smooth with no bunching at zero. The smoothness of the distribution combined with its centering at 30% suggests it is unlikely the government is optimally targeting market values with guidance values. To the right of the bunching region, the market value distribution resembles a right-shifted version of the $\frac{r-c}{c}$ distribution; this is consistent with substantial tax inaccuracy via under-payment of taxes relative to market values. The fraction of transactions with $\frac{p-c}{c} < 0$ is 14.5%, consistent with some tax-payers either over-paying taxes or facing costly audit procedures to certify their lower tax values.

Figure 2b contrasts the sum of tax base values ($\max[r, c]$) (blue triangles) with the sum of estimates of total market values within $\frac{r-c}{c}$ bins (red left-arrows). The aggregation of rupee amounts within bins in this plot helps to alleviate concerns that classical measurement error in our market value proxy p could make the p distribution in Figure 2a resemble a smoothed version of the reported value distribution. Intuitively, if the market value proxy is measured with such symmetric error, aggregating many transactions within each bin should take care of this issue, clearing out the noise in the proxy and facilitating more accurate comparison with aggregated reported values. Consistent with agents under-reporting taxes, we find that the aggregate amounts measured by the market value proxy consistently lie beneath aggregate reported values, with a particularly large gap located precisely in the bin where reported values bunch at the guidance value. Appendix Section C.1 develops this argument further with simulations.

Finally, Figure 2c plots the percentage under-reporting rate (blue triangles) estimated within $\frac{r-c}{c}$ bins. This figure is constructed by taking the percentage difference between the total tax base value within a bin and our estimate of total market value within a bin (i.e., the percentage difference between the blue triangles and the red left-arrows in Figure 2b). There is a spike in the estimated under-reporting rate for transactions reporting exactly at c; moreover transactions with $\frac{r-c}{c} < 0$ are still primarily under-paying, with this behavior declining to zero when $\frac{r-c}{c}$ is approximately 0.5. We discuss below the ways in which this apparent decline could be an artifact of one type of measurement error in transaction value proxies; see Appendix C for further discussion and plots with a measurement error correction.

We note, however, that these comparisons using binned distributions do not directly reveal transaction-specific over- and under-payment of taxes relative to market values, because the transactions in a given bin across the two series are not necessarily the same.[29] Appendix Figure C13a measures tax inaccuracy by plotting the total

[29] Some transactions that report $\frac{r-c}{c} < 0$ might have $\frac{p-c}{c} > 0$, and therefore still under-pay. And in

amount of inaccuracy of tax payments relative to our proxy measure of accurate tax payments τp (in reality, an accurate tax payment is τm). The blue points in the figure show the absolute value of the difference between an accurate tax payment and the actual tax payment $\tau(p - \max[r, c])$, divided by estimated total market value within the bin (p). These averages are calculated within 2% bins of the $\frac{r-c}{c}$ ratio. Overall, we see the aggregate amount of tax inaccuracy hovers between 1% and 1.5%; a large amount relative to the desired tax rate τ which is 5% in our context.

The green and red curves in Appendix Figure C13a decompose the total inaccuracy into the portion arising from under-payment and over-payment of taxes, again relative to our estimate of total market values p. Moving left to right, we first observe that even when households report $\frac{r-c}{c} < 0$, we estimate that they are still primarily under-paying taxes. This suggests that guidance values are set quite low to prevent people from being forced to over-pay; in the $\frac{r-c}{c} < 0$ region over-payments represent only 0.25% of the p reported in that bin. In the region of $\frac{r-c}{c} \approx 0$, we see a spike in under-payments relative to our estimated market values. This is consistent with $r = c$ bunchers being particularly likely to under-pay.

Mortgages can provide an important incentive to curb misreporting, as we discuss in Section 2.1. In our setting, mortgages also have to be reported to the registrar. To make progress, we match transactions with mortgages, which are separately reported to the registrar. We are able to match roughly 31,000 reported transactions to a mortgage.[30] Of these 31,000 mortgage-matched transactions, a smaller number, 8,913, also match to a Propequity value p.

To more fully utilize the matched data, we therefore take a different approach, plotting how bunching behavior varies with loan-to-value ratios for the full set of 31,000 matched mortgage transactions. Figure A8 shows that transactions with progressively higher loan-to-value ratios tend to exhibit less bunching and that this relationship is monotonic. This finding is consistent with our argument that incentives to relax credit constraints (higher reported values lead to the possibility of greater mortgage loans) cut against incentives for tax evasion. The magnitudes are sizeable—low loan-to-value loans are approximately 10 percentage points more likely to bunch than transactions associated with high loan-to-value mortgages. As we do not utilize exogenous variation in credit constraints here, we note that this result could also be driven by a negative correlation between preferences for tax evasion and agents' credit constraints.

the case of measurement error, a transaction that reports $r > c$ might also have $r > p$, which we would measure as an over-payment even though the buyer reports $r > c$.

[30] We assume that any transaction that is unmatched to a mortgage transaction has a loan-to-value (LTV) ratio of zero, since this biases us against finding any differences between zero LTV and positive LTV transaction samples (this mechanically makes the zero LTV sample more similar to the positive LTV sample, because some of the zero LTV sample likely have a positive but unmeasured LTV).

4.1 Accounting for Measurement Error in the Market Value Proxy

We seriously consider the possibility that the market value proxy p that we source is measured with error, and adopt several approaches to deal with this. First, we directly use a subset of the data that has plausibly less measurement error where we only use the p for any given transaction when there is an exact match of property locations across the two data sources (reported values to the registrar and our source for the proxy for market value). We find similar results using this exact match sample (see Appendix Figure C15).

Second, we also conduct a higher frequency analysis of transaction and reporting behavior. These tests are discussed in detail in Appendix Section C; we summarize the main findings here. We consider two alternative explanations: (1) the government optimally sets the guidance value c to reflect true market prices m; and (2) buyers and sellers use c as a reference point when negotiating transaction prices m. Under either scenario, the discrepancy between the true market price and the potentially noisy proxy p could lead us to mis-estimate the extent of over- or under-payment. If guidance values c closely tracked market prices m for these reasons, we would expect the most pronounced bunching to occur immediately after c is updated—since those values are least stale relative to current market conditions (given that c is only updated annually). In contrast, if bunching is primarily driven by under-reporting, we would expect it to peak just before a guidance value update, when under-reporting at $r = c$ is most profitable due to a lower c. Figure 3a shows that bunching is highest in the days just prior to guidance value revisions. Regression evidence confirms that under-reporting is significantly higher in months preceding c changes (Appendix Table A1). Second, we find that transaction timestamps are clearly consistent with agents back-dating transactions to exploit lower c values (Figure 3b). These patterns are difficult to reconcile with alternative explanations and instead support strategic timing to minimize tax burdens.

Finally, as we show in Appendix Figure C13a, the estimated fraction of over-payment increases with $\frac{r-c}{c}$ to the right of zero. This estimated over-payment comes from an increasing fraction of transactions reporting $r > p$ as we move right in $\frac{r-c}{c}$ space. It is possible that this pattern is due to real reporting behavior; for example buyers might over-report to reduce future capital gains taxes by establishing a higher than market cost basis, but it is also potentially driven by measurement error. In Appendix Section C.1.1, as a robustness check, we develop a measurement error correction to the proxy p to deal with this issue, and re-estimate our structural model. Even with this correction, we find similar results for our main parameter of interest (the government's willingness to pay to avoid tax inaccuracy).

5 Structural Estimation

Our goal in this section is to provide the first estimate of a government's revealed preference for tax accuracy versus tax revenue (i.e. the parameters ψ_{under} and ψ_{over} in equation (8)). We first estimate the parameters that govern households' optimal reporting and evasion behavior. With those parameters in hand, we go on to estimate the ψ values that gives us the closest match to the observed enforcement policy (guidance values, in our context).

5.1 Parameters Governing Individual Reporting Behavior r^*

Four estimated parameters govern optimal reporting/evasion behavior in our structural model of tax reporting. An informal discussion of identification of these parameters follows. First, $\pi_{r\geq c}$ is the household's perceived audit probability when they choose to report more than the guidance value $r > c$.[31] A higher audit probability, ceterus paribus, delivers a higher reporting amount. The level of the audit probability also determines the size of the discrete change to this probability that the household faces when moving r from just below to just above c. More specifically, since $\pi_{r\geq c} < 1$, households have a strong incentive to either bunch their reported r values exactly at c, or slightly to the right of c ("diffuse-bunching"), to avoid a certain audit (as well as associated transactions costs/penalties) from reporting $r < c$.

The second parameter, ρ, governs household beliefs about whether evasion is detected conditional on being audited. Lower ρ values—leading to a higher expected penalty—lead some buyers to report higher than others conditional on the same m, thus capturing this source of heterogeneity in the data.

The parameter t determines the transaction costs faced by households from being audited. Higher t values lead households to avoid reporting $r < c$, as in such cases, households are audited pay these transaction costs with certainty. The parameter β affects how much t scales by the difference between r and c—capturing the extent to which households with similar m values report different r values both just above and just below c.

ϵ_i determines household i's belief about the government's audit behavior. Positive values of ϵ_i mean households will bunch to the right of the guidance value c, believing that they can avoid a certain audit by reporting slightly higher values than the guidance value—the distribution of ϵ_i thus captures the diffuse bunching observed in the data.

[31] Recall, we fix the household's audit probability at 1 if they choose to report $r < c$ as this automatically triggers an audit.

5.1.1 Details of Structural Estimation

We parameterize heterogeneity in the structural parameters using single-parameter distributions for the sake of parsimony. We draw $\rho \sim \chi^2(K)$, searching for K over a grid spanning $[0, 4]$. Given that $\pi_{r\geq c} \in (0.25, 1)$ (as discussed earlier), we search over this support for $\pi_{r\geq c}$.[32] $\epsilon_i \sim \chi^2(E)$, where we search a grid from $[0, 2]$ for E. This value for ϵ_i is expressed in terms of c, i.e., if $\epsilon_i = 0.5$, then the individual believes the jump from $\pi_{r\geq c}$ to $\pi_{r<c}$ is at $1.5 \times c$. t is drawn from a $\chi^2(T)$ searched for over the grid $[0, 1]$, as it is expressed as a rate relative to c (i.e. a $t = .2$ implies a transaction cost of audit equal to $.2c$). We set β to 1, meaning that the transactions cost of appeal scales one for one with the gap $c - r$, thus eliminating the need to estimate an additional parameter.

We observe data on reported (r_i), guidance (c_i) and (estimated) market values (m_i). The transaction tax in Mumbai $\tau = .05$ and the penalty multiplier $n = 4$. We estimate model parameters using the count and rupee value bunching distributions in Figure 2 panels (a) and (b) as our targets, deploying a minimum distance estimator (MDE).[33] More specifically, we jointly grid search over all parameters, and simulate from the resulting distributions for each parameter choice, thus generating optimal reporting amounts r_i^* from the model, which we use to find the optimal parameter set to minimize the following loss function:

$$\sqrt{\left(\frac{1}{K}\sum_{k=1}^{K}\omega_k\frac{(r_k - r_k^*)^2}{\bar{r_k}}\right) + \left(\frac{1}{K}\sum_{k=1}^{K}\omega_k\frac{(v_k - v_k^*)^2}{\bar{v_k}}\right)}, \tag{9}$$

where k are the 2% $\frac{(r-c)}{c}$ bins from -1 to 1 (a total of 99 bins), r_k is the total observed reported counts, r_k^* the total model-implied reported counts within each bin, v_k the total observed reported value, and v_k^* the model-implied reported values within each bin for each combination of parameters. ω_k is set such that for a given bin midpoint mid_k, we have $\omega_k = 1 - |mid_k|$. For example, if a 2% $\frac{r-c}{c}$ bin has a midpoint of 0.06, the associated weight is 0.94.[34] Since we equally weight two different sets of moments with different units (counts and rupee-values), we normalize count and rupee errors by the average bin count $\bar{r_k}$ and average bin value $\bar{v_k}$ observed in the data respectively.

We find that the optimal value of $\pi_{r\geq c} = 0.30$, $K = 1.02$, $T = 0.33$, and $E = 0.80$,

[32] Recall that the lower bound for $\pi_{r\geq c}$ is $\frac{1}{n}$, as otherwise, the penalty is irrelevant in expected value terms, leading to the degenerate case of the household always reporting c .

[33] Since the distributions that we use as empirical targets are not raw moments of the data, we are conducting minimum distance estimation, as opposed to a method of moments. In the terminology of MDE (e.g., Cameron and Trivedi (2005), Chapter 6.7), the count and value distributions are reduced-form parameters that we use to infer the deep parameters of the model.

[34] This is similar to a simple form of weighting by the fraction of total observations in each bin, broadly equivalent to inverse variance weighting, accounting for the higher density of observations close to c.

characterizing the distributions of ρ, t, and ϵ_i respectively. The perceived probability of audit $\pi_{r \geq c}$ is 30%, substantially lower than one, capturing the fact that in the data, we observe many fewer transactions to the immediate left of c relative to immediately to the right of c in the data. This parameter value implies that the probability of audit (and therefore the need to pay transaction costs t) changes substantially (by 0.70) around c, so agents have a strong incentive to report just at or above c versus below. $\rho = 1.02$ implies a generally low belief of detection probabilities conditional on audit; with $\rho = 1.02$ the bunchers in our analysis (who under-report by approximately 25%, see Figure A2) believe that there is a 24% detection probability conditional on audit.[35] This parameter results primarily from the substantial amount of under-reporting we observe in the data, both at c and away from the guidance value.

The estimate of $E = 0.8$ implies that approximately half of all buyers believe the probability of audit is 1 even if they report 35% above c. This high-level of E matches the data pattern of substantial diffuse bunching around the guidance value c. However, given the low estimated values of ρ, this estimate of E can be reinterpreted—these households essentially believe that even a sure audit is unlikely to discover the full amount of under-reporting.

Figure 2 panels (a) and (b) show the best model fit as green dots for both counts (panel (a)) and rupee values (panel (b). Panel (c) of this figure also shows how the model (once again in green dots) matches an untargeted moment, namely, the estimated under-reporting rate from the data. Overall the model fit to the data matches the main features of the reporting data. We therefore turn next to using these estimated parameters governing household reporting behavior to ascertain the government's revealed preference for tax accuracy.

5.2 Estimating the Government's Preference for Tax Accuracy

To aid intuition about how we identify ψ values in our framework, Figure 4 presents comparative statics of how total welfare varies as the government adjusts different instruments, with a particular focus on c, the guidance value. In each plot, we use the observed distribution of market values (i.e., p) together with the parameters governing taxpayer reporting (estimated in section (5.1.1)) to simulate how the distribution of r^* varies with each choice of government instruments (such as c and n). The simulated distribution of r^* is then simply aggregated according to the government's objective function (equation (8)) to obtain total welfare, which is the y-axis of all plots in Figure 4. The x-axis of each of these plots is the government guidance value c; for ease

[35] $(1 - .75)^{1.02} = 0.24$.

of interpretation, we normalize the empirically observed government guidance value to 1, meaning that movements along the x-axis of these figures can be interpreted as percentage increases (to the right of 1) and decreases (to the left of 1) from the empirically observed value. Put differently, in these figures, movements in the single policy parameter c proportionally scales up or down all guidance values across the city of Mumbai.

We first consider the baseline case when the inaccuracy term in the government welfare function is $\psi = \psi_{over} = \psi_{under}$, i.e., symmetric concern about over- and underpayment. Figure 4a shows how welfare varies with the guidance value c in this case; each curve in this plot corresponds to a different level of the inaccuracy weight parameter ψ. When ψ increases, the government cares relatively more about inaccuracy of both kinds, giving the tax authority greater incentive to lower c. However, beyond a certain high level of ψ, c^* converges to a constant, meaning that when over- and underpayment are treated symmetrically, the optimal policy becomes relatively insensitive to higher penalties on inaccuracy.[36]

Figure 4b shows how welfare varies with adjustments to the penalty multiplier n. Increases in n for a given ψ allow the government to set lower levels of c, since individuals generally report higher r values when misreporting penalties increase.

Figures 4c and 4d relax the assumption that the government cares symmetrically about tax inaccuracy arising from people who under-pay or over-pay. We assume in turn that the tax authorities care only about under-payers, i.e., $\psi_{over} = 0$, in panel (c); and that they only care about over-payers, i.e., $\psi_{under} = 0$, in panel (d). Each curve in these plots represents a different level of the non-zero ψ value (ψ_{under} and ψ_{over}, respectively), showing how it changes the role of c. When the authorities only care about under-payers in panel (c), they need to set c very high to simultaneously fulfill their objectives to maximize revenues and drive out under-payers. However, when the tax authority cares only about over-payers in panel (d), c falls rapidly to values much lower than under the symmetric case.

Structurally estimating the government's weight on inaccuracy ψ: The analyses in Figure 4a, 4c and 4d suggest that existing policy—in which guidance values are generally lower than market transaction values on average—are consistent with the government placing higher priority on avoiding inaccurate overpayments as in Figure 4d. This priority is also consistent with the observation that there are often complaints from households when they face guidance values c greater than their m values. Fi-

[36] While it appears as though the empirically observed guidance value ($c = 1$) can be rationalized in this case with relatively high levels of ψ, it is worth reiterating that this plot is illustrative, and unrealistically treats the entire city of Mumbai as a single subzone driven by a single policy instrument.

nally, the government's revenue-maximization objective inherently prefers higher reporting amounts, diminishing the need to impose a separate penalty on under-payers. We therefore adopt this specification with $\psi_{under} = 0$ and for our remaining structural analyses we estimate ψ_{over}, which we denote ψ for simplicity.

We further assume that ψ does not differ across regions, i.e., it is the same for the whole city of Mumbai, though we do allow model-implied c to vary across regions of the city as observed in the data. That is, we set as our target moments the observed guidance values c_{st} for sub-zones s in years t in our data. The algorithm proceeds heuristically as follows. We first pick a ψ value. Given this ψ we proceed to estimate the vector of guidance values c^* that maximizes the government's objective function (which includes both revenue maximization and tax accuracy maximization as specified in equation (8). We repeat this process for different values of ψ, finally picking the ψ^* that produces a c^* that is closest to the observed guidance values c in the data according to the loss-function:

$$\min_{\psi} \sqrt{\sum (C^*_{st}(\psi) - C_{st})^2 \omega_{st}}, \tag{10}$$

where ω_{st} represents weights for each subzone s and time t—these place greater weight on sub-zone $\times$ year clusters with more observations, a simple form of inverse-variance weighting.

In structurally estimating ψ^*, we assume that the government does not observe the true transaction values m for all individual transactions. However, we do assume that the government does observe the distribution of m within each subzone. We believe this assumption is reasonable because if the government had access to the true values of m, guidance values would be unnecessary. Furthermore, this assumption realistically implies that the government cannot anticipate which specific properties will be transacted—only the distribution from which these transactions are drawn.

As suggested by the shape of the p distribution observed in Figure 2a, we assume that m, when normalized by c, follows a normal distribution[37]. Since the welfare weights in equation (7) require the transaction amounts (rather than amounts normalized by the guidance value), we simulate m, $\frac{m}{c}$ observations jointly from a multivariate normal distribution. For each subzone, these simulated values of $\frac{m}{c}$ and m form the basis for estimating ψ, and the parameters estimated in Section 5.1 produce the corresponding reported values r^*. As discussed, this approach enables us to compute welfare for any given combination ψ and guidance value vector c. By maximizing welfare within each subzone, we obtain a unique optimal guidance value vector c^* correspond-

[37] Since both the mean and standard deviation need to be estimated, for accuracy we limit our analysis to subzones with more than 50 observations in a given year.

ing to each specified level of ψ. Finally, we recover the government's implied value of ψ^* by minimizing the loss function in equation 10.[38]

We structurally estimate $\psi^* = ₹2.12$; this implies the government is willing to give up ₹2.12 per ₹1 of over-payment. We interpret this as a substantial magnitude; a government that only focuses on maximizing tax revenue would have a ψ of zero. Intuitively, guidance values that aimed to just maximize tax revenue would be substantially higher, but the constraint of having a preference for not over-taxing some agents too much leads the government to essentially "under-enforce" tax policy.[39]

5.3 Validation: Predicting Guidance Values Out-of-Sample

Our estimation, which yields parameters governing taxpayer reporting behavior, as well as an estimate of ψ^*, relies on several assumptions. To validate whether we have appropriately characterized these economic primitives, we conduct an out-of-sample forecasting exercise. We check whether the optimal guidance values calculated based on our structural model of government objectives and reporting behavior predict changes to observed guidance values.

To conduct this exercise, we first calculate optimal guidance values for each year based on our structural model of reporting, but estimate the parameters of the structural model of reporting using data from all other years. For instance, for the year 2016, we structurally estimate ψ^* and the associated vector c^* (for all sub-zones in the data) using data from all years *except for 2016*.

We compare the optimal out-of-sample c^* to the observed guidance values c in Figure 5a, which plots a histogram of the ratio $\frac{c^*}{c}$ computed for all years in our dataset. The modal observation is roughly 1.3, meaning that our estimates of c^* broadly suggest that policymakers have room to raise guidance values for some sub-zones, assuming that we have accurately captured their objectives. The histogram also reveals several subzones with $c^* = 0$. These areas commonly exhibit $p < c$ as well as substantial cross-transaction variance in p. Given the value of ψ^*, the government's preference for fairness outweighs its revenue objective in these cases, and leads to the model predicting $c^* = 0$ in such cases. If geography is a contributing factor to the high variance in p, this suggests that finer subzone partitioning may be justified.

We next use the distance between the model-implied c^* and observed c for each sub-zone to forecast revisions to c over the next year (for example, estimated 2016 dif-

[38] As discussed earlier, we assume that the extensive margin elasticity to c is zero—in support of this assumption, Appendix Section D presents evidence that in the case of Mumbai, revisions to c do not result in any meaningful changes to the total volume of transactions.

[39] Implementing the measurement error correction discussed in C, we estimate $\psi = ₹2.3$. The government's preference for tax accuracy is similar with or without this correction.

ferences between model-implied c^* and observed c are used to forecast 2017 revisions to c out-of-sample in the example above). This checks whether we can predict up or down revisions to the guidance value in regions where the optimal guidance value is lower or higher than the prevailing guidance value. Figure 5b presents the average change to guidance value predicted by the model alongside the realized changes. The top panel presents the (binned, sorted) out-of-sample prediction using the model, and the bottom-panel shows the out-of-sample realization that ensues within each predicted change bin. The figure shows that both upward and downward revisions are well-predicted by deviations between model-implied c^* and actual c. While these directional forecasts are broadly in line with the data, the magnitude of observed revisions are small in magnitude relative to our model—consistent with other factors at play to explain quantitative responses than in our simple structural model.[40] Nevertheless, this evidence serves as useful out-of-sample validation of our model.

6 Conclusion

We develop a new framework for optimal taxation and enforcement when policymakers care about both welfare maximization and tax accuracy. This framework accommodates the widespread sentiment that there is value in reducing tax noncompliance separate from the revenues raised, and a corresponding cost to over-collecting revenues beyond what is legally owed. Our model formalizes these sentiments by adding a priority for tax accuracy in the policymaker's objective function.

We apply this framework to the empirical setting of property transaction taxes in Mumbai with self-reported transaction values. We structurally recover the degree of misreporting, and the elasticity of misreporting and transaction volume to enforcement, based on the degree of bunching in reported valuations around government-assessed guidance values which serve as a floor on the tax base. Existing policy suggests a strong preference against tax inaccuracy—and especially against tax overcollection—on the part of policymakers, which we validate using an out-of-sample forecasting exercise.

References

Agarwal, Sumit, Keyang Li, Yu Qin, Jing Wu and Jubo Yan. 2020. "Tax evasion, capital gains taxes, and the housing market." *Journal of Public Economics* 188:104222.

[40] Appendix Table A2 confirms these visual results in a regression.

Agarwal, Sumit, Weida Kuang, Long Wang and Yang Yang. 2020. "The Role of Agents in Fraudulent Activities: Evidence from the Housing Market in China.".

Allingham, Michael G. and Agnar Sandmo. 1972. "Income tax evasion: A theoretical analysis." *Journal of Public Economics* 1(3):323–338.

Amornsiripanitch, Natee. 2020. "Why Are Residential Property Tax Rates Regressive?" *Available at SSRN 3729072* .

Anagol, Santosh, Allan Davids, Benjamin B. Lockwood and Tarun Ramadorai. 2024. Diffuse Bunching with Frictions: Theory and Estimation. Technical report National Bureau of Economic Research.

Anagol, Santosh, Fernando V Ferreira and Jonah M Rexer. 2021. Estimating the Economic Value of Zoning Reform. Technical report National Bureau of Economic Research.

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx and Tarun Ramadorai. 2022. "Reference dependence in the housing market." *American Economic Review* 112(10):3398–3440.

Avenancio-León, Carlos F and Troup Howard. 2022. "The assessment gap: Racial inequalities in property taxation." *Quarterly Journal of Economics* 137(3):1383–1434.

Badarinza, Cristian, Tarun Ramadorai, Juhana Siljander and Jagdish Tripathy. 2024. "Behavioral Lock-In: Housing Market Taxation with Reference Dependent Agents." *Available at SSRN 4693047* .

Best, Michael Carlos and Henrik Jacobsen Kleven. 2018. "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK." *Review of Economic Studies* 85(1):157–193.

Bukovina, Jaroslav, Tomáš Lichard, Ján Palguta and Branislav Žúdel. 2025. "Corporate Minimum Tax and the Elasticity of Taxable Income: Evidence From Administrative Tax Records." *American Economic Journal: Economic Policy* 17(2):358–387.

Cameron, Colin A and Pravin K Trivedi. 2005. *Microeconometrics: Theory and Applications*. Cambridge University Press.

Casaburi, Lorenzo and Ugo Troiano. 2016. "Ghost-house busters: The electoral response to a large anti–tax evasion program." *Quarterly Journal of Economics* 131(1):273–314.

Chang, Yun-chien. 2012. "Self-Assessment of Takings Compensation: An Empirical Study." *The Journal of Law, Economics, & Organization* 28(2):265–285.

Comptroller and Auditor General of India. 2016. Report of the Comptroller and Auditor General of India on Revenue Sector for Year Ended 21 March 2015. Technical report.

Dachis, Ben, Gilles Duranton and Matthew A Turner. 2012. "The effects of land transfer taxes on real estate markets: evidence from a natural experiment in Toronto." *Journal of Economic Geography* 12(2):327–354.

Dzansi, James, Anders Jensen, David Lagakos and Henry Telli. 2024. Technology and tax capacity: evidence from local governments in Ghana. Technical report National Bureau of Economic Research.

Fan, Ziying, Xiaxin Wang and Hang Zhang. 2022. "Understanding Misreporting: New Evidence from Responses to a Housing Transaction Tax.".

Garmaise, Mark J and Tobias J Moskowitz. 2004. "Confronting information asymmetries: Evidence from real estate markets." *Review of Financial Studies* 17(2):405–437.

Gechter, Michael and N Tsivanidis. 2020. "Spatial spillovers from urban renewal: evidence from the Mumbai Mills Redevelopment." *Penn State University (mimeo)* .

Genesove, David and Christopher Mayer. 2001. "Loss aversion and seller behavior: Evidence from the housing market." *Quarterly journal of economics* 116(4):1233–1260.

Gopalakrishnan, S. and A. Das-Gupta. 1986. "Stamp Duties and Registration Fees in West Bengal." Unpublished Manuscript.

Harari, Mariaflavia and Maisy Wong. 2018. Slum upgrading and long-run urban development: Evidence from Indonesia. In *2018 Meeting Papers. Society for Economic Dynamics*.

Harberger, Arnold C. 1965. *Issues of tax reform for Latin America*. PAU, Department of Economic Affairs.

Keen, Michael and Joel Slemrod. 2017. "Optimal tax administration." *Journal of Public Economics* 152:133–142.

Khan, Adnan Q, Asim I Khwaja and Benjamin A Olken. 2016. "Tax farming redux: Experimental evidence on performance pay for tax collectors." *Quarterly Journal of Economics* 131(1):219–271.

Kopczuk, Wojciech and David Munroe. 2015. "Mansion tax: The effect of transfer taxes on the residential real estate market." *American economic Journal: economic policy* 7(2):214–57.

Mayle, Peter. 2000. *A year in Provence*. Penguin UK.

Mayshar, Joram. 1991. "Taxation with costly administration." *The Scandinavian Journal of Economics* pp. 75–88.

Mirrlees, James A. 1971. "An exploration in the theory of optimum income taxation." *The review of economic studies* 38(2):175–208.

Mirrlees, James A. 1986. "The theory of optimal taxation." *Handbook of mathematical economics* 3:1197–1249.

National Institute of Public Finance and Policy. 1995. "Black Money in the Real Estate Sector: A Study.".

Okunogbe, Oyebola and Fabrizio Santoro. 2022. "The promise and limitations of information technology for tax mobilisation.".

Panchapagesan, V. and N. Karthik. 2017. "Revenue Neutral Approach to Lower Stamp Duty and Registration Charges for Affordable Housing." Unpublished Manuscript.

Pomeranz, Dina and José Vila-Belda. 2019. "Taking state-capacity research to the field: Insights from collaborations with tax authorities." *Annual Review of Economics* 11:755–781.

Posner, Eric A and E Glen Weyl. 2019. *Radical Markets*. Princeton University Press.

Ramsey, Frank P. 1927. "A Contribution to the Theory of Taxation." *The economic journal* 37(145):47–61.

Regan, Tanner. 2023. "The Economic Consequences of Statistical Property Valuation: Evidence from Kampala, Uganda." Abstract.

Rocha, Roberto, Thiago Scot and Javier Feinmann. 2023. "Property Taxes in Sao Paulo.".

Saez, Emmanuel and Stefanie Stantcheva. 2016. "Generalized social marginal welfare weights for optimal tax theory." *American Economic Review* 106(01):24–45.

Sandmo, Agnar. 1981. "Income tax evasion, labour supply, and the equity—efficiency tradeoff." *Journal of public economics* 16(3):265–288.

Sher, Itai. 2024. "Generalized social marginal welfare weights imply inconsistent comparisons of tax policies." *American Economic Review* 114(11):3551–3577.

Slemrod, Joel, Caroline Weber and Hui Shan. 2017. "The behavioral response to housing transfer taxes: Evidence from a notched change in DC policy." *Journal of Urban Economics* 100:137–153.

Sood, Aradhya. 2020. Land Market Frictions in Developing Countries: Evidence from Manufacturing Firms in India. Technical report Mimeo. https://aradhyasood. github. io/Sood_Land_Frictions_India1. pdf.

Srinivasan, T.N. 1973. "Tax evasion: A model." *Journal of Public Economics* 2(4):339–346.

Tandon, N.K. 1987. "Under-reporting of Real Estate Transactions in India." Unpublished Manuscript.

Tsivanidis, Nick. 2019. "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio." *Unpublished manuscript* .

Yitzhaki, Shlomo. 1979. "A note on optimal taxation and administrative costs." *The American Economic Review* 69(3):475–480.

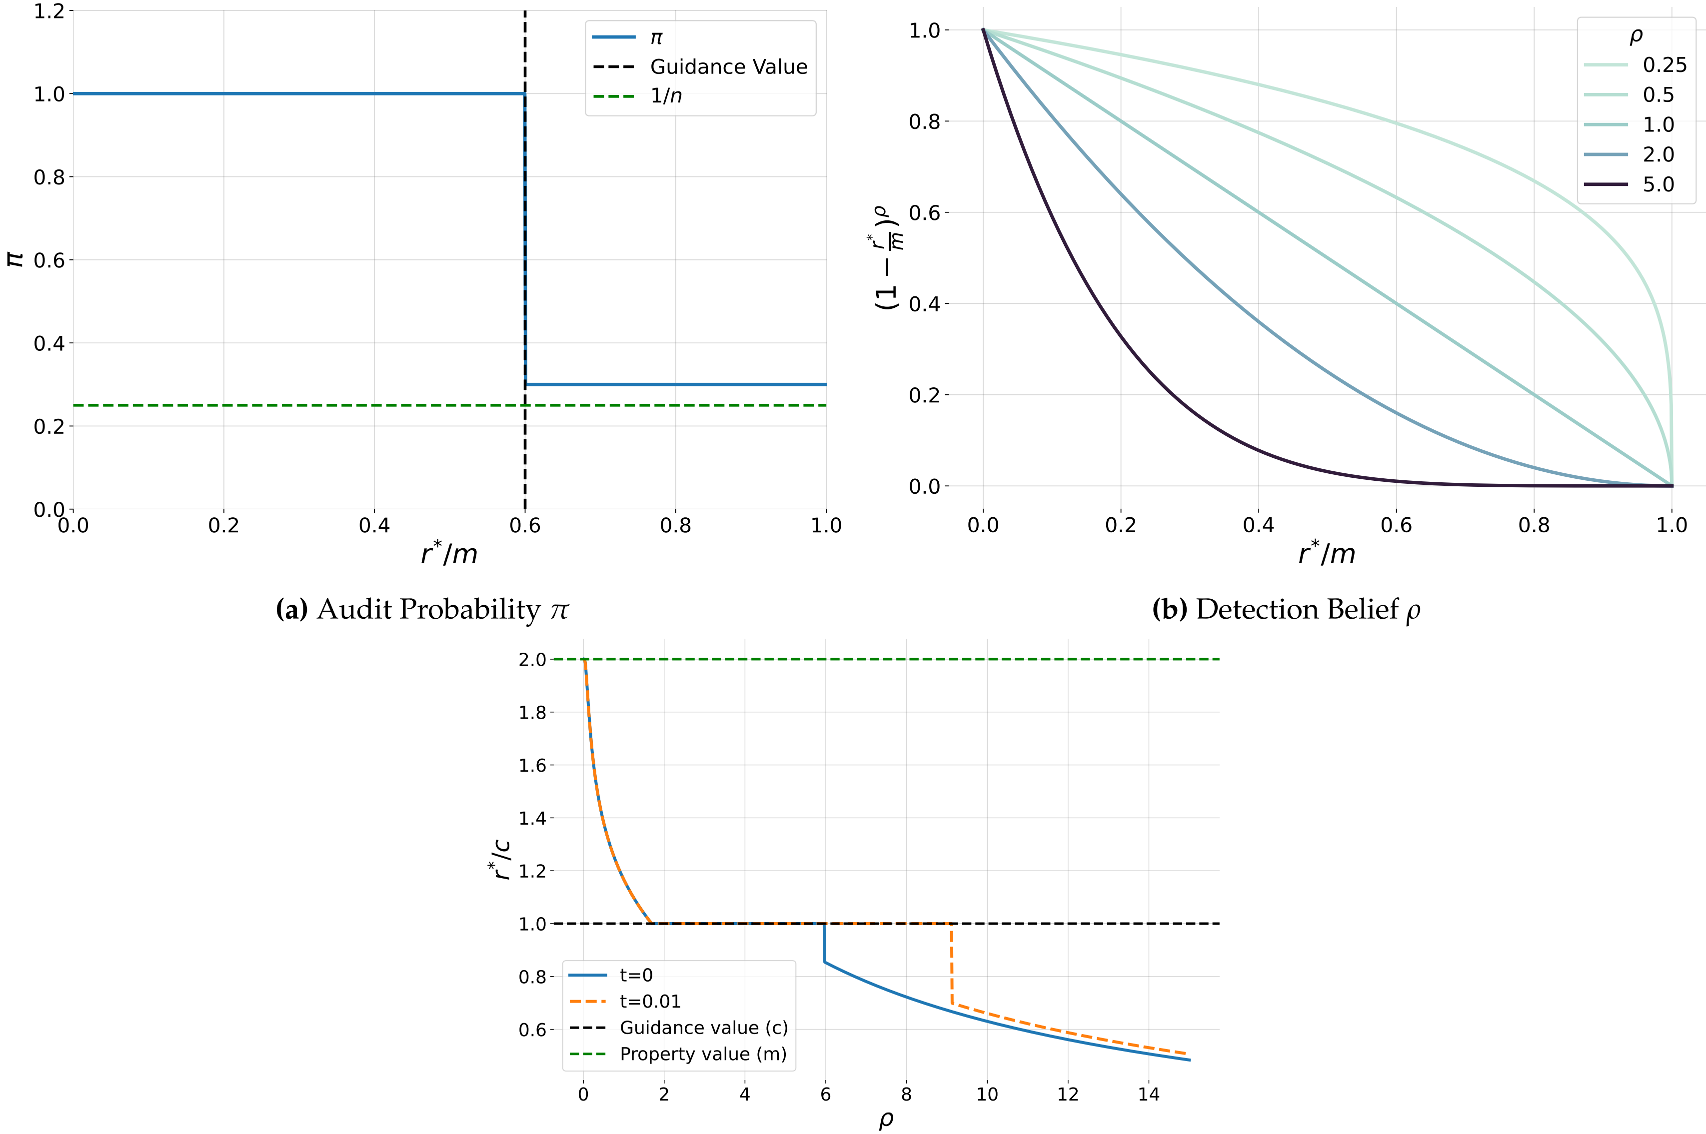

(a) Audit Probability π

(b) Detection Belief ρ

(c) Reported Values

Figure 1
Belief Parameters in the Model

Panel (a) presents the audit probability π, and Panel (b) presents the detection belief parameter ρ. Panel (c) presents reported values, relative to c, for various values of ρ, and in the presence of transaction costs, t.

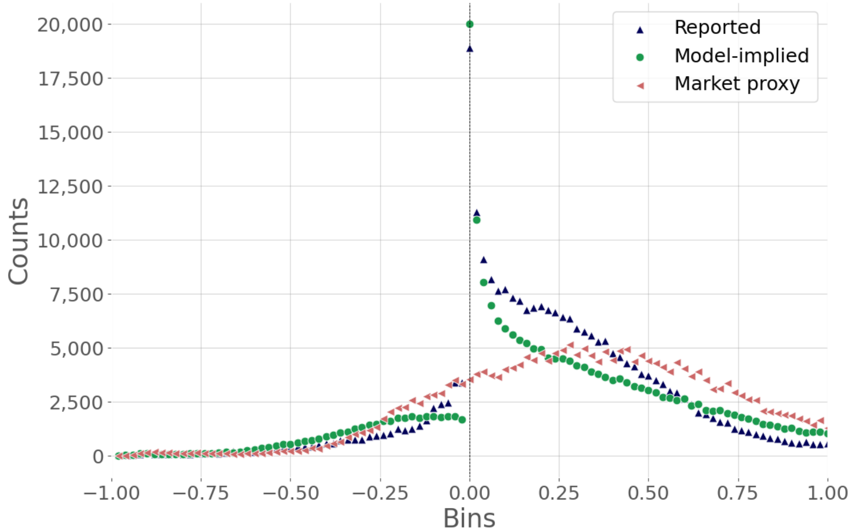

(a) Count Distribution

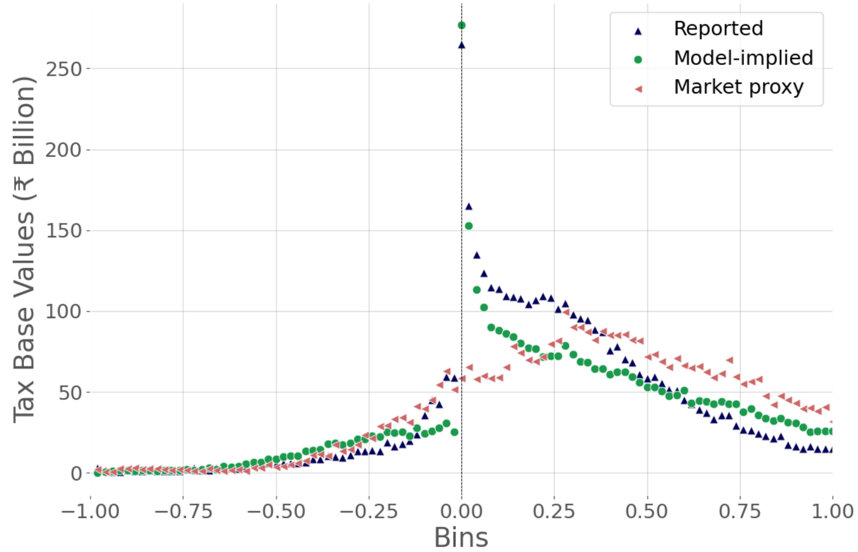

(b) Value Distribution

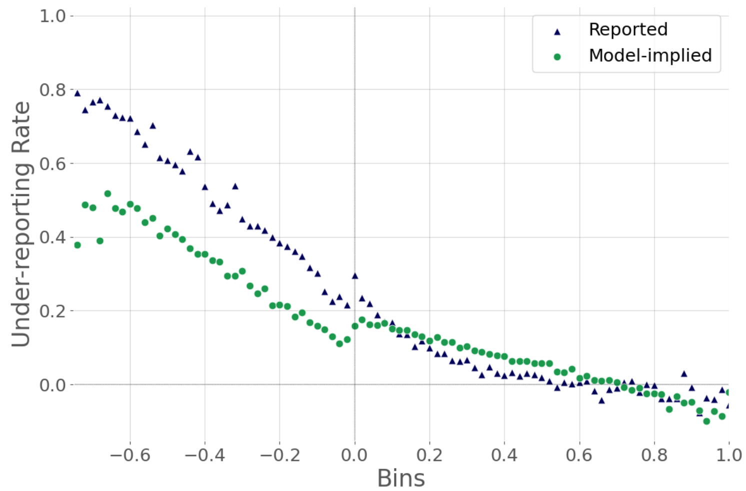

(c) Under-reporting Rate

Figure 2
Raw Data and Structural Model Fits

Panel (a) plots the distribution of reported values r, market value proxies p, and our model's predicted distribution of r_{model} around the guidance value c. Panel (b) plots the corresponding tax base, defined as $\max[x, c]$ with $x \in \{r, p, r_{model}\}$, in billions of rupees. Panel (c) presents the model fit to the untargeted under-reporting rate of reported values r. All figures use $\frac{x-c}{c}$ on the x-axis.

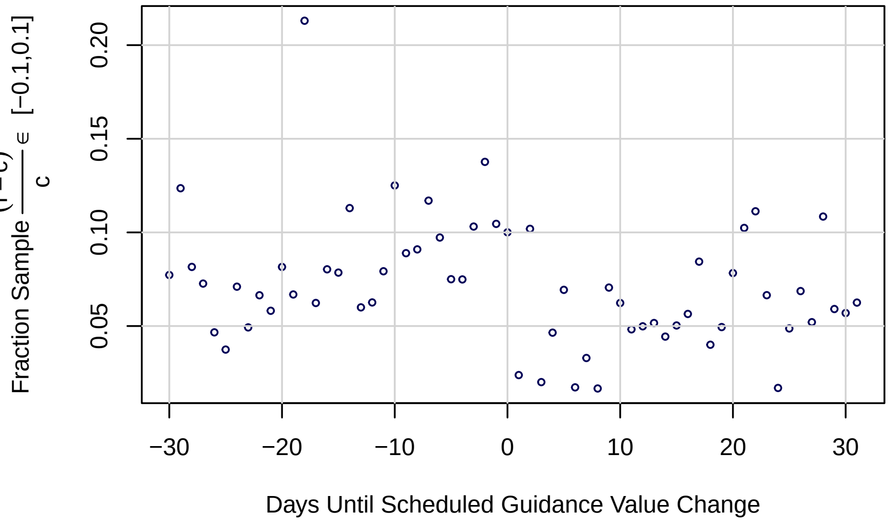

(a) Fraction Bunching

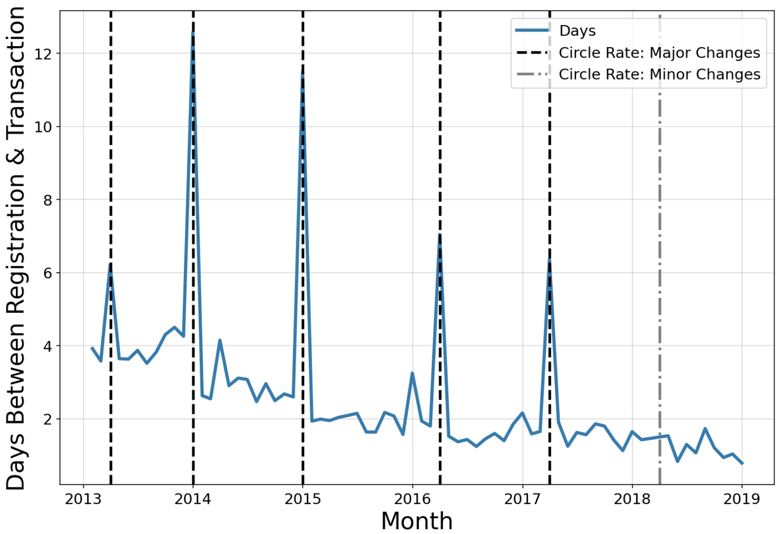

(b) Average Days Transaction Registered Before Completed

Figure 3
Agreement Date Backdating and Transaction Counts Before Guidance Value Changes

Panel (a) plots the fraction of transactions that bunch within 1% of the guidance value in the days before and after guidance value increases. Panel (b) plots the agreement month on the x-axis and the average number of days between the registration and transaction dates on the y-axis.

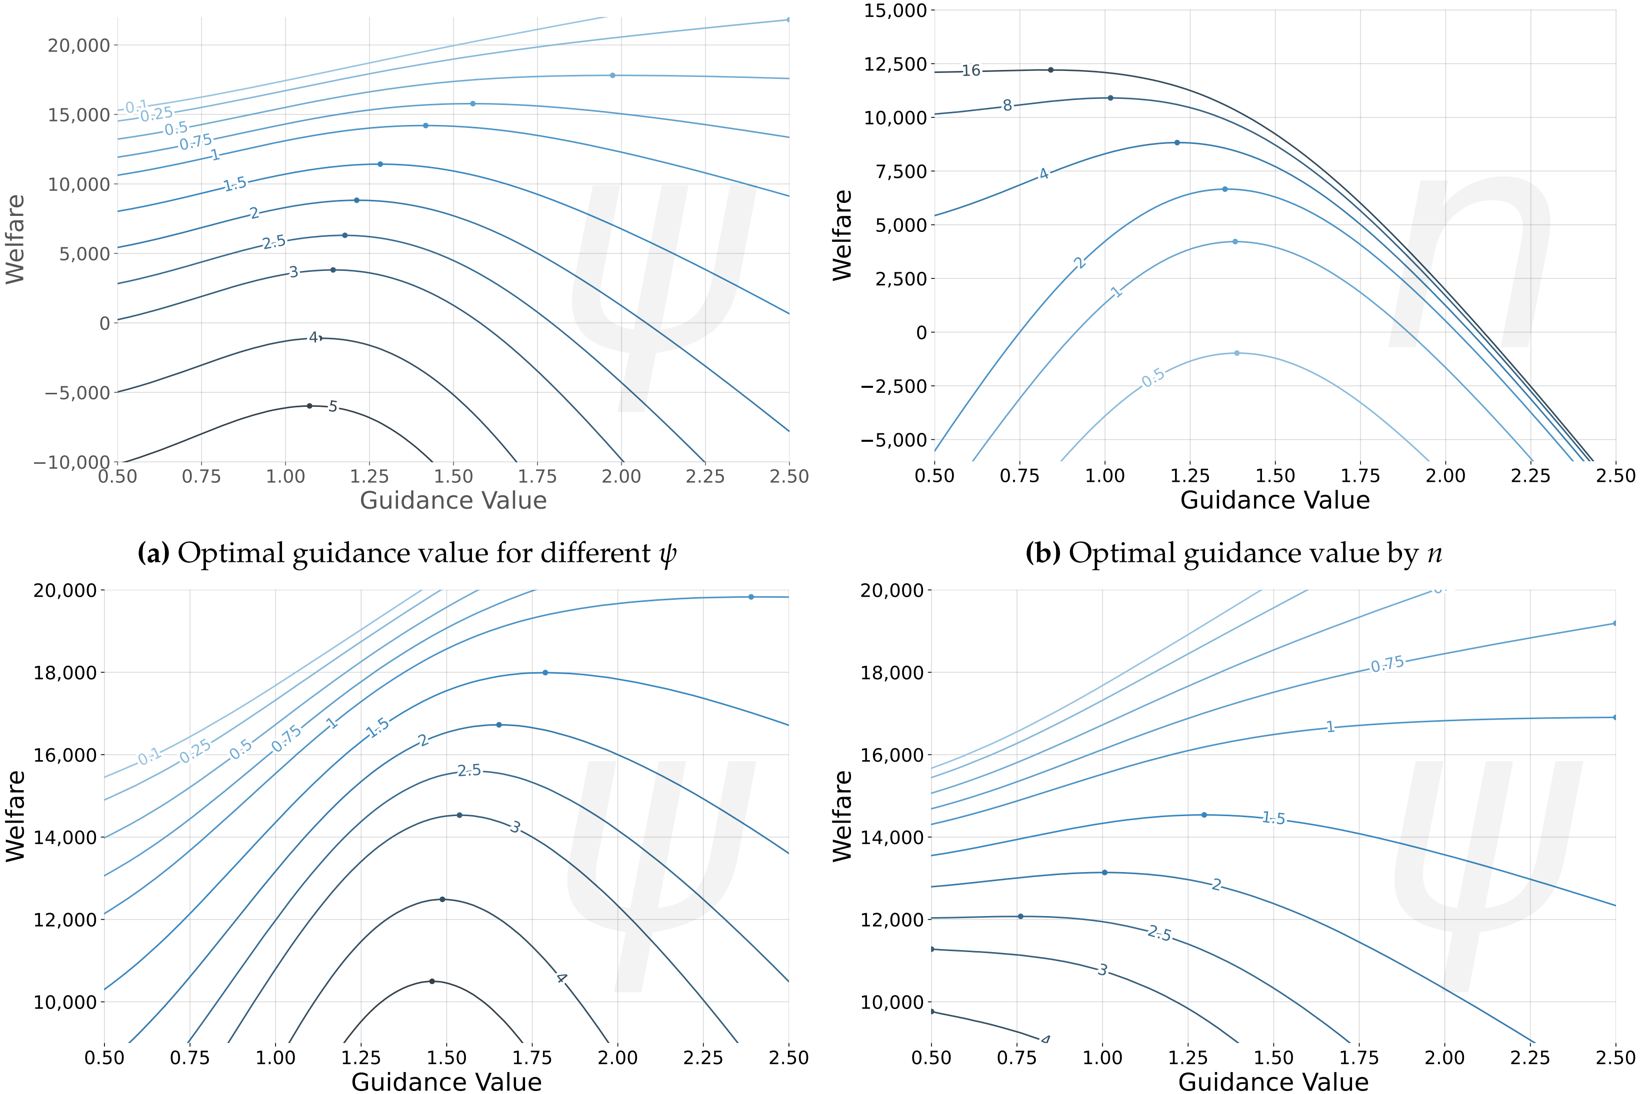

(a) Optimal guidance value for different ψ

(b) Optimal guidance value by n

(c) Optimal guidance value when only under-payers matter

(d) Optimal guidance value when only over-payers matter

Figure 4

Comparative Statics: Optimal Government Policy

Panel (a) presents the comparative statics for optimal guidance value varying ψ, (b) for various values of the penalty parameter n, Panel (c) when only under-payers matter (therefore finding the guidance value should be as high as possible, to avoid under-payers) and Panel (d) when only over-payers matter.

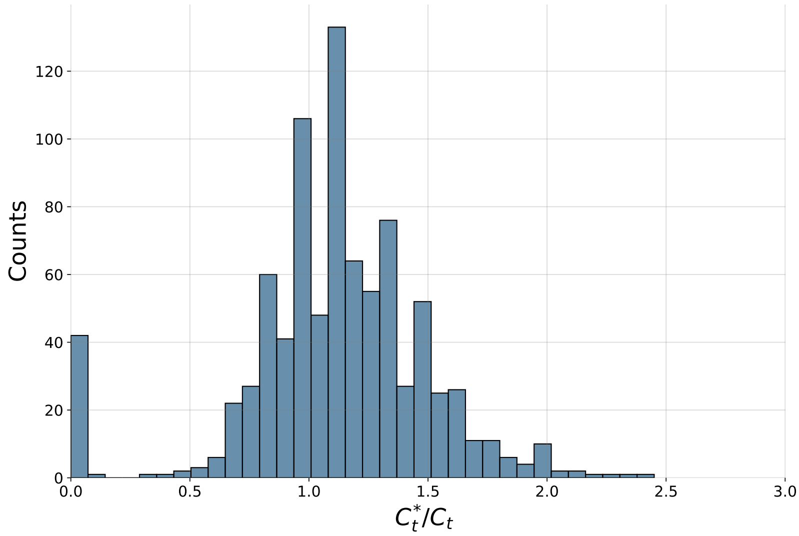

(a) c^*/c: Histogram

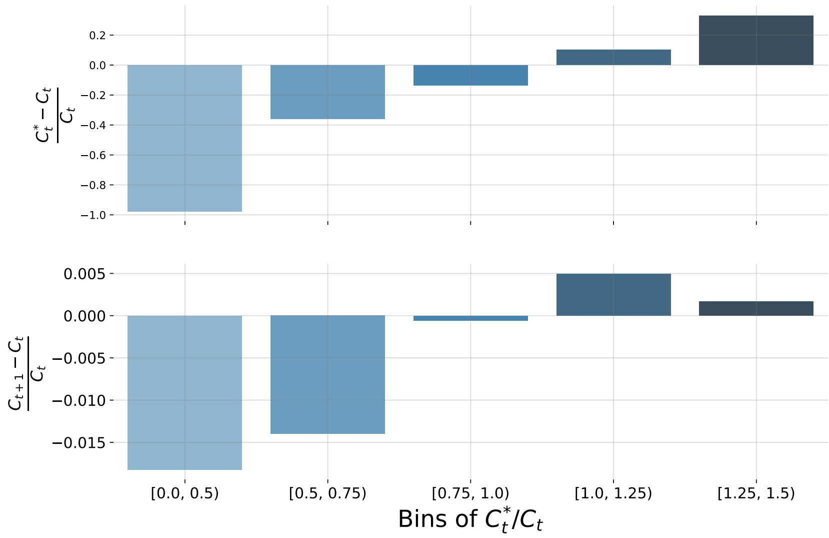

(b) Predicting Revisions in c by c^*/c

Figure 5
Predicting Revisions in c

Panel (a) presents the histogram of distance between c^* and c, while Panel (b) presents the average distance within each c^* distance bin (top sub-panel) and the revision in the bottom sub-panel.

	Reported Value '000s USD		Guidance Value '000s USD		Propequity Value '000s USD		Primary Transaction = 1		Area (sq M)		No. Obs.
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median	
2013	275.66	173.01	208.71	146.04	326.32	205.23	0.66	1	85.55	76.58	13,648
2014	313.27	195.57	234.15	158.61	362.44	225.50	0.69	1	88.75	75.14	17,213
2015	319.13	203.41	251.66	175.52	362.99	243.08	0.70	1	83.16	72.46	20,615
2016	315.70	205.50	258.22	175.88	363.41	241.19	0.70	1	79.13	69.90	23,803
2017	337.22	220.00	289.13	194.27	386.11	258.64	0.74	1	78.00	68.82	31,104
2018	324.03	212.86	264.82	175.48	367.51	243.07	0.76	1	73.53	65.04	38,228
2019	315.59	215.04	254.19	171.50	360.73	245.33	0.73	1	71.02	62.97	30,602
2020	320.36	210.00	266.78	172.50	371.71	241.19	0.72	1	71.35	62.34	30,289
2021	334.36	212.52	274.85	177.39	374.59	246.06	0.68	1	72.21	62.15	49,663
2022	324.45	220.00	257.77	176.44	355.16	240.91	0.63	1	68.14	60.57	5,449
Total	321.77	210.00	261.86	174.42	367.29	242.02	0.71	1	76.06	66.28	260,614

Table 1

Summary Statistics on Transactions

The table reports summary statistics for the set of transactions that is either matched to the same project from Propequity or to the nearest Propequity project. A primary transaction is one where the housing unit is sold by a real estate developer.

Online Appendix: Optimal Tax Policy with Misreporting: Theory, and Evidence from Real Estate

Santosh Anagol   Vimal Balasubramaniam
Benjamin B. Lockwood   Tarun Ramadorai   Antoine Uettwiller

August 22, 2025

A Supplementary Tables & Figures Cited in Main Text

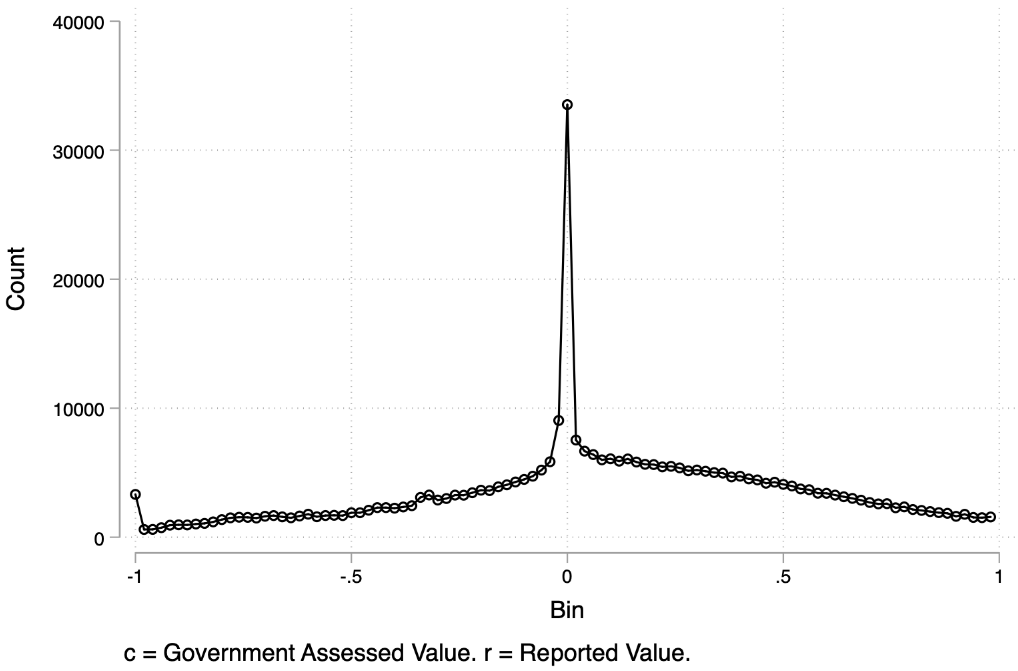

Figure A1
Bunching Estimates for Sao Paulo, Brazil

This figure replicates Rocha, Scot and Feinmann (2023) who show similar bunching patterns using the ITBI municipal transaction tax on properties in Sao Paulo, Brazil. The transaction tax rate is 3%, and it is charged on the higher of the buyer's reported value or a guidance value. The line shows the distribution of reported values across 2% reported value bins, where a reported value bin is measured as a deviation from the guidance guidance value.

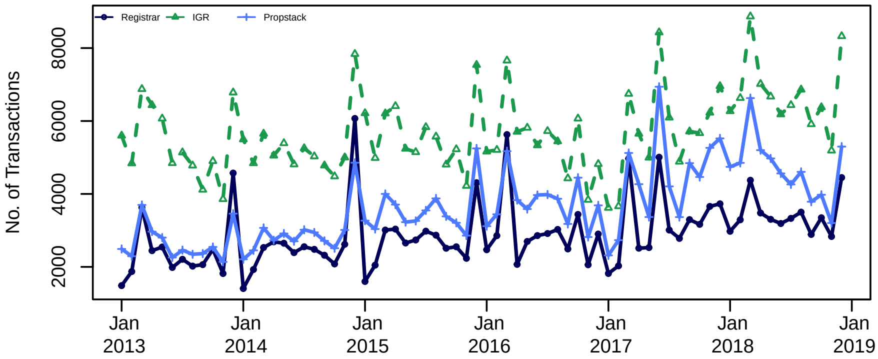

(a) Transactions

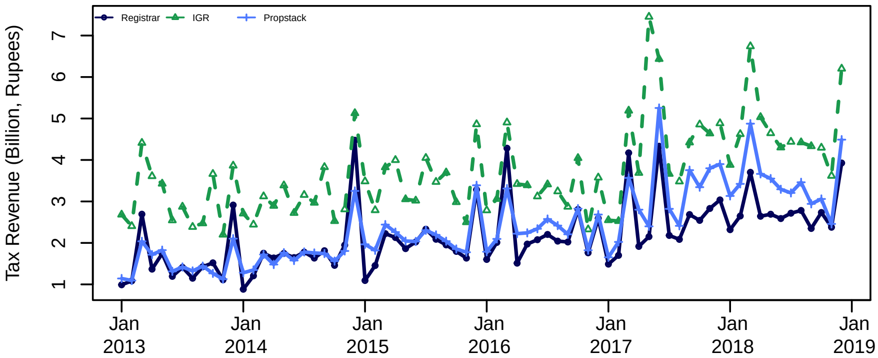

(b) Tax Revenue

Figure A2
Sample Comparison to Aggregate Tax Revenues

This figure plots the monthly time series of the total number of transactions in panel (a) and the total tax revenue from these transactions in panel (b). The blue line with circles plot the numbers obtained from aggregating the extracted Registrar data, the green triangle is the sum reported by the Inspector General of Registrations for a region that Mumbai and Mumbai suburban areas belong to, that is larger than our sample, and the light blue line with "+" plots the aggregated information from Propstack analytics. The overlapping data sample period ends in January 2019, although our full sample is between 2013–2022.

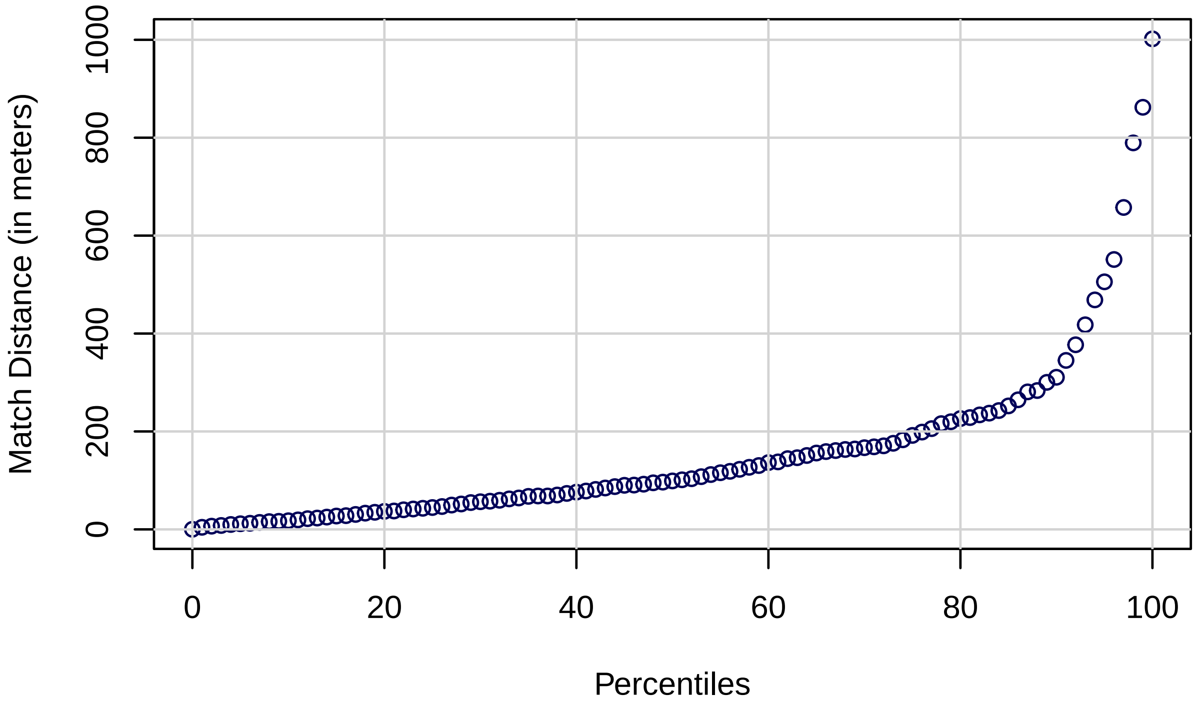

Figure A3
Propstack-Propequity Match Quality

This figure plots the empirical distribution of the match distance (in meters) between a transaction in our Propstack data (source of administrative data on reported and guidance values) and the match from our Propequity data (source of estimated price data). 80% of the transactions matched to Propequity price data are within 200 meters, and 95% of the transactions are matched to Propequity transactions within 500 meters.

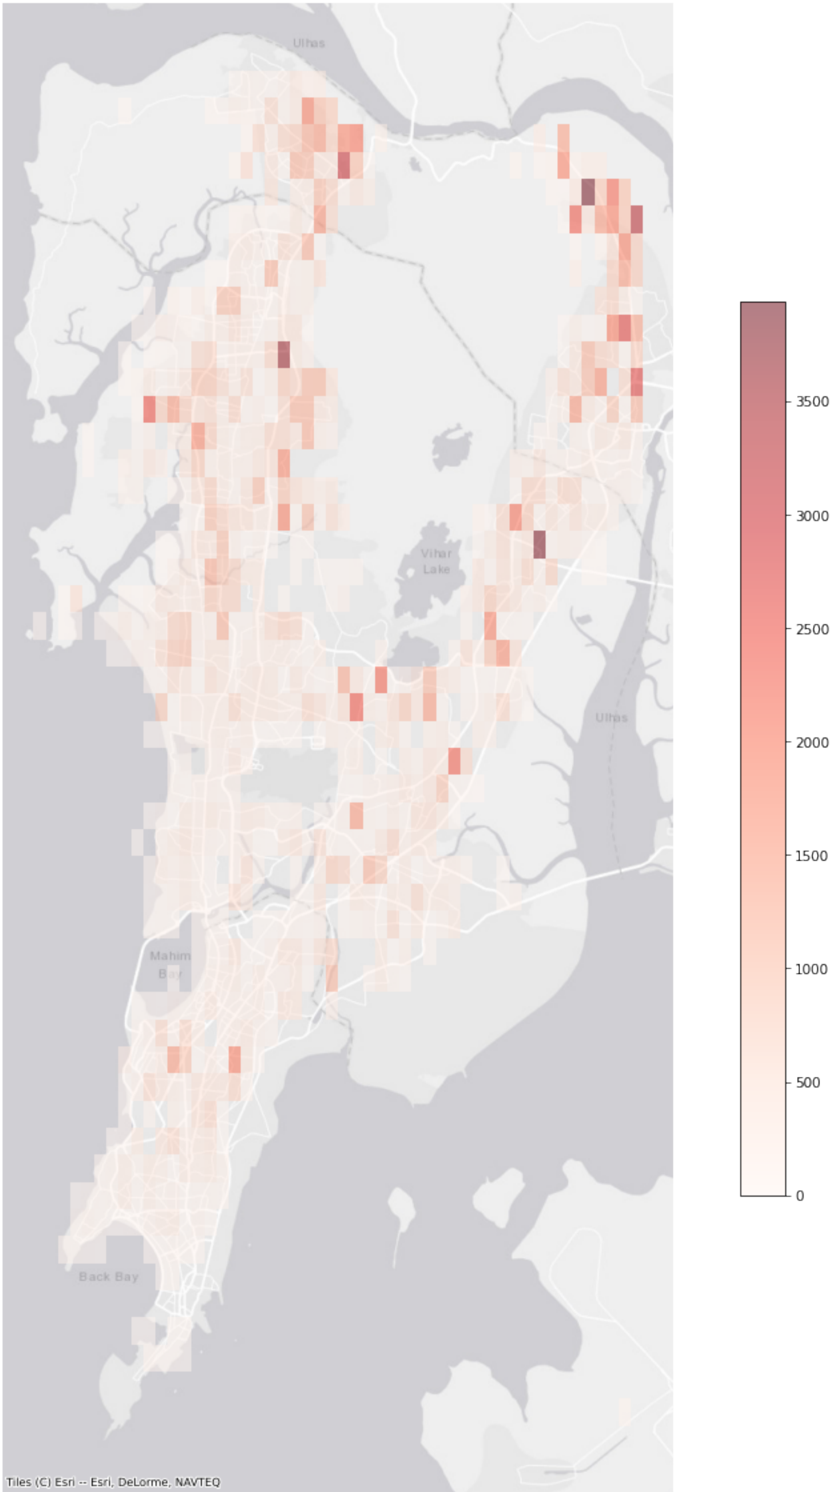

Figure A4
Heatmap of Transactions in our Final Sample

This heatmap presents the spatial distribution of the final set of transactions in our sample in Mumbai and Mumbai Suburban regions between 2013–2022.

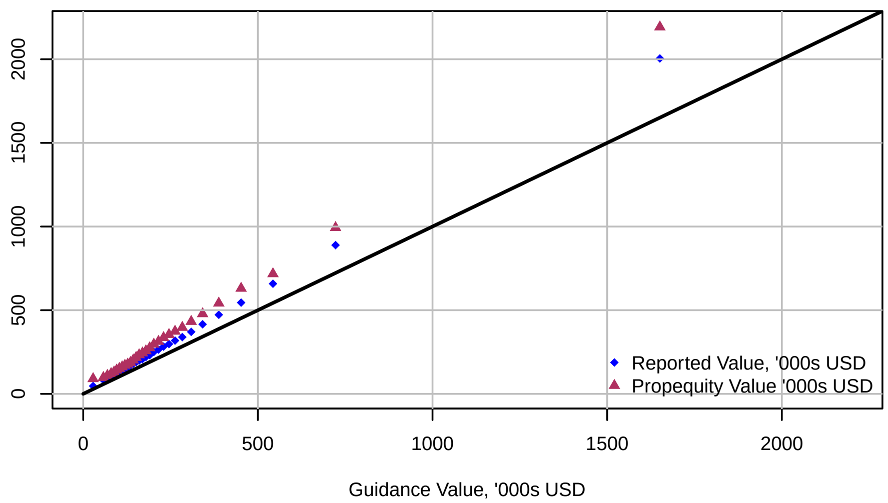

Figure A5
Correlation between Reported Value, Propequity (Estimated Market) Values, and Guidance Values

This figure presents a binned scatterplot of the average reported values (blue diamonds) and estimated market values from the Propequity data (maroon triangles) within guidance value bins. The black line is the 45-degree line.

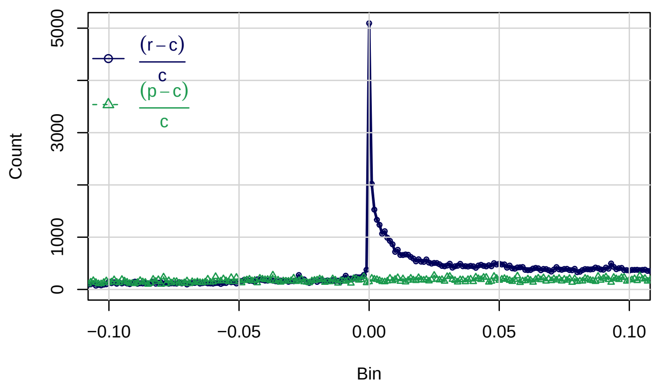

Figure A6
Bunching of Reported and Propequity Values Around Circle Values in 0.1% Bins

The blue line/circle shows the distribution of reported values across 0.1% reported value bins, where a reported value bin is measured as a deviation from the guidance value. The green line/triangle shows the distribution of our noisily measured estimate of the market price (the Propequity values) for the same underlying set of transactions reported in the blue line/circle.

				2BHK COST SHEET-
Saleable Area		**1237**	**1242**	**12**
Basic Rate		6099	6099	
Basic Cost		7544463	7574958	
Car Parking Charges (1 car park slot)		400000	400000	
Preferential Location Charges (Rate per sqft)		50	50	
Preferential Location Charges Cost		61850	62100	
Basic Unit Value		**8006313**	**8037058**	**8061**
				Additi
Corpus Fund Rs 40/- per sq.ft		49480	49680	
Infrastructure charges Rs 75/- per sq.ft		92775	93150	
Maintenance Charges for 12 months @ Rs. 3.50/sq.ft		51954	52164	
CMWSSB & TNEB Charges		100000	100000	
Documentation Charges		25000	25000	
Club Amenities charges		50000	50000	
GST@18%		57551	57657	
Total Additional Charges Including GST@18%		**426760**	**427651**	**428**
				PAYMEN
Application Amount		300000	300000	
Allotment stage Payment (To be paid on or before15th day from Booking Date)	**10%**	500631	503706	
1st Quarterly Instalment - 11th Quarterly Instalment (To be paid on or before 45th day from Booking Date)	**85%**	6805366	6831499	
12 Quarterly Installment upon Intimation on Possession (Inclusive of Additional charges)	**5%**	827076	829503	
Total Payable		**8,433,073**	**8,464,709**	**8,**

Figure A7
Example Price Sheet from Propequity

This extract presents the detailed breakdown of the costs covered by our data provider. Our estimate of the market value includes both the property purchase cost and other ancillary services.

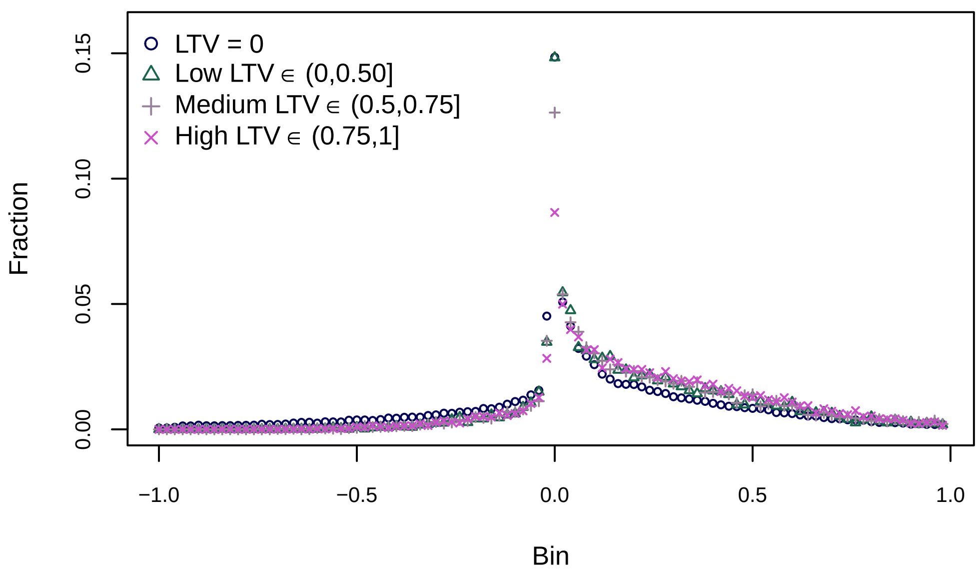

Figure A8
Bunching and External Financing Constraints

This figure presents the fraction of observed transactions with mortgages within 2% bins, with the bin around zero ranging from −1% to +1%, by various loan-to-value bins.

Table A1
Under-reporting Before Circle Rate Changes

Dep Var: Under-reporting Rate	(1)	(2)	(3)	(4)
Month Before Policy Change	0.061***	0.054***	0.064***	0.068***
	(0.000)	(0.000)	(0.000)	(0.000)
Mean Dep Var.	0.06	0.06	0.06	0.06
Intercept	Yes	Yes	Yes	Yes
Time-trend	No	Yes	Yes	Yes
Month of year FE	No	No	Yes	Yes
Year FE	No	No	No	Yes
No. Obs.	260,614	260,614	260,614	260,614

The table reports the regression results estimating the average under-reporting rate the month before circle rate changes.

Table A2
Predicting Revisions in Guidance Value using Structural Estimates

	Pct. increase in c	Pct. increase in c
Percent Distance to c^*	0.040***	0.021**
	(0.002)	(0.006)
Year Fixed Effects	No	Yes
Subzone Fixed Effects	No	Yes
Number of Years	9.0	9.0
Number of Subzones	197.0	197.0
Obs	868	868

Standard errors in parentheses. * $p < .05$, ** $p < .01$, *** $p < .001$

The table reports the regression results predicting the percentage change in c, the guidance value, in the next year, using the percentage deviations from the optimal guidance value c^* implied by the structural estimates of our model.

B Data

B.1 Transactions Data

Our primary dataset on reported values and guidance values comes from Propstack Analytics and is described in the main text. This section describes the underlying documents from which the Propstack Analytics data is ultimately sourced.

The main underlying data source on transactions is the publicly available individual property transaction reports released by the Office of the Inspector General of

Registration and Controller of Stamps (IGR), Department of Revenue, Government of Mahrashtra, India. This state apparatus plays an important role in collecting state government revenues from across the state using various fiscal instruments available in the state government's toolkit. The state is split into 8 regional divisions, and we obtain data for the Mumbai regional division, which comprises Mumbai City and Mumbai Suburban districts. Our study area currently covers 437 square kilometers out of the 6,640 square kilometers Mumbai Metropolitan Region. We currently focus on this region because we can reliably obtain transaction data that can be mapped to geo-spatial information relevant for our study.

The *eSearch* facility set up by the IGR enables access to transaction-level data for all properties transacted in Greater Mumbai. Every transaction report is in Marathi, the most commonly spoken language in Maharashtra. Figure B9 presents an example of the original document downloaded from the IGR *eSearch* facility. Figure B10 presents the transaction report translated into English using Google's translation services. The details available in each transaction report provide a consistent information set for all real-estate transactions for Greater Mumbai. This information set also serves as the basis for the government to make policy decisions on real estate transaction taxes.

Each transaction report obtained from the *eSearch* facility begins with a document number and the name of the registrar office (the local IGR office for a region). The more substantive information is in the form of a table starting with the name of the local village where the property is located[41]. The first row of the table in Figure B10 lists the type of transaction. All real-estate transactions in Maharashtra are classified as "Agreement", "Agreement to Sale", "Sale deed" and "Transfer Deed" types. We filter all downloaded transaction reports to these deed types to form our core data set.

The second row lists the reported price at which the transaction took place. In this case, the reported transaction price is ₹7,500,000. The third row lists the price as per the government-issued guidance value, known as the *policy circle rate* that is determined annually by a legally predetermined process. The policy circle rate determines the floor price at which the government will deem this property to be sold for taxa-

[41] Historically, the Mumbai region was formed of seven islands or fishing villages, which then expanded rapidly over time. The village tag to geographies is more of an artefact of historical documentation than a reference to the economic or social conditions of different regions in Mumbai.

tion purposes. The value of this property, according to government-determined circle rates, is ₹4,434,062. Row 12 provides the computed stamp tax paid on this transaction of ₹375,000, determined as the prevailing stamp tax rate, in this case 5%, on the reported transaction value. The circle rate plays an important role in that it sets the lower bound for the stamp tax revenue generated for a property of this type. In the event this property's reported transaction price is below ₹4,434,062, the stamp duty payable will be 5% of this guidance value, after which the law facilitates a process by which the related parties can file for revision. The fourth row of this table provides the property address, and other measurement details in terms of the area of the house, the land registry survey number, and other information relevant for determining the circle rate. The fifth row of this table lists the property area, and the next few rows list details of the two parties to the transaction. Row 9 reports the transaction date for the document, and Row 10 reports the actual date of formal registration for the sale. These two dates can be different as the law allows for a grace period of 3 months from the actual transaction date, during which time they are legally bound to register the sale with the registrar. The last row provides data on the registration fee paid, which is capped at ₹30,000 or 1% of the reported value, whichever is lower.

We validate the coverage of our transaction reports data from the IGR *eSearch* facility by matching the total real-estate stamp duty number of documents filed and revenue generated in each year from our transactions data to the official aggregate numbers. Figure A2 presents this comparison. The top panel of the figure shows the official number of documents filed with the IGR in each month in green against the number of documents in our transaction data (orange for the registrar data we manually sourced for matching to mortgages and blue for Propstack). The time series are very highly correlated. The bottom panel of Figure A2 shows the official aggregate tax revenue collected from stamp duty in each month against the aggregate stamp revenue from our transaction data. The total revenue figures from both sources include stamp revenues and the registration fees for all transactions in a given month. Once again, the time series are very highly correlated, especially in the second half of our sample period

Although we capture a majority of the transactions in Greater Mumbai, the differ-

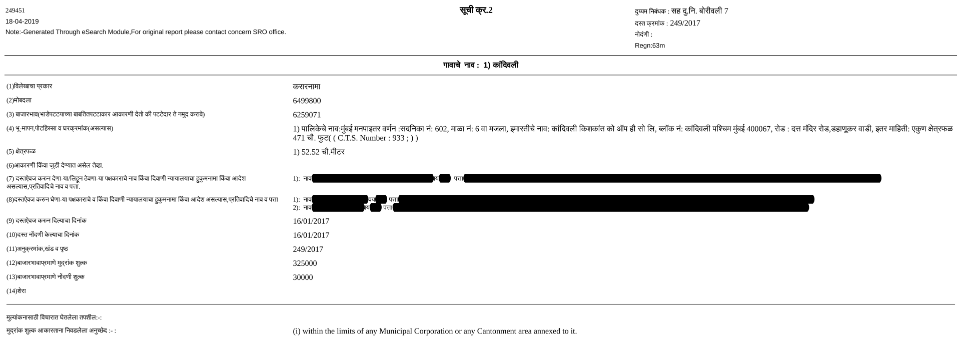

249451
18-04-2019
Note:-Generated Through eSearch Module,For original report please contact concern SRO office.

सूची क्र.2

दुय्यम निबंधक : सह दु.नि. बोरीवली 7
दस्त क्रमांक : 249/2017
नोदंणी :
Regn:63m

गावाचे नाव : 1) कांदिवली

(1)विलेखाचा प्रकार	करारनामा
(2)मोबदला	6499800
(3) बाजारभाव(भाडेपटटयाच्या बाबतितपटटाकार आकारणी देतो की पटटेदार ते नमुद करावे)	6259071
(4) भू-मापन,पोटहिस्सा व घरक्रमांक(असल्यास)	1) पालिकेचे नाव:मुंबई मनपाइतर वर्णन :सदनिका नं: 602, माळा नं: 6 वा मजला, इमारतीचे नाव: कांदिवली किशकांत को ऑप हौ सो लि, ब्लॉक नं: कांदिवली पश्चिम मुंबई 400067, रोड : दत्त मंदिर रोड,डहाणूकर वाडी, इतर माहिती: एकुण क्षेत्रफळ 471 चौ. फुट((C.T.S. Number : 933 ;))
(5) क्षेत्रफळ	1) 52.52 चौ.मीटर
(6)आकारणी किंवा जुडी देण्यात असेल तेव्हा.	
(7) दस्तऐवज करुन देणा-या/लिहून ठेवणा-या पक्षकाराचे नाव किंवा दिवाणी न्यायालयाचा हुकुमनामा किंवा आदेश असल्यास,प्रतिवादिचे नाव व पत्ता.	1): नाव[redacted]वय[redacted] पत्ता[redacted]
(8)दस्तऐवज करुन घेणा-या पक्षकाराचे व किंवा दिवाणी न्यायालयाचा हुकुमनामा किंवा आदेश असल्यास,प्रतिवादिचे नाव व पत्ता	1): नाव[redacted]वय[redacted] पत्ता[redacted] 2): नाव[redacted]वय[redacted] पत्ता[redacted]
(9) दस्तऐवज करुन दिल्याचा दिनांक	16/01/2017
(10)दस्त नोंदणी केल्याचा दिनांक	16/01/2017
(11)अनुक्रमांक,खंड व पृष्ठ	249/2017
(12)बाजारभावाप्रमाणे मुद्रांक शुल्क	325000
(13)बाजारभावाप्रमाणे नोंदणी शुल्क	30000
(14)शेरा	

मुल्यांकनासाठी विचारात घेतलेला तपशील:-:
मुद्रांक शुल्क आकारताना निवडलेला अनुच्छेद :- : (i) within the limits of any Municipal Corporation or any Cantonment area annexed to it.

Figure B9
Original Transaction Document: An Example

This figure presents the original transaction document obtained from the IGR *eSearch* facility `https://freesearchigrservice.maharashtra.gov.in`.

ences between our aggregate revenue numbers and the official figures arise from two reasons. The official figures for Mumbai also include a suburban area of Navi Mumbai, which we do not include in our sample. Moreover, we count revenues in the month the transaction was registered, and this may not necessarily be the same as the official approach, especially for transactions that may be executed at the end of the month, but are fulfilled in the early periods of the following month.

B.2 Circle Rate Scheduled Changes

Circle rates are set at the sub-zone level, a geographic area of approximately .67 square kilometers on average. Table B3 presents the summary statistics on the circle rate variation in Greater Mumbai for our sample period. At the start of our sample period, we have 727 sub-zones, which increase to 747 sub-zones in 2015, and then stabilize at 734 for the remainder of our sample period.[42] In the early years of our sample, nearly all sub-zones underwent changes in circle rates. The average change in each year varies from 0% in 2018 to 14.4% from the previous year in 2015 (Column 3).

[42] New sub-zones are formed by either dividing existing sub-zones into multiple new ones, or by fusing different parts of multiple sub-zones to form new ones. We keep track of all of the changes in the geo-spatial files, thus identifying which regions form to create the new sub-zones, and the old sub-zones they belonged to.

249451
18-04-2019
Note: -Generated Through eSearch Module

List no. 2

Sub Registrar with duni Borivali 7
Document number: 249/2017
Note:
Regn: 63m

Name of the village: 1) Kandivali

(1) Type of document	Agreement
(2) Reward	6499800
(3) Quotes (The leaseholder leases the details of the rent that the sergeant should specify)	6259071
(4) Land measuring, portals and home number (if any)	1) Name of the corporation: Mumbai Manipatayar Description: House No. 602, Malala No. 6th Floor, Name of the building: Kandivali Kishakant to Op Hau Soil, Block No. Kandivli West Mumbai 400067, Road: Datta Temple Road, Dahanukar Wadi, Others Information: total area 471 square Foot (CTS Number: 933;))
(5) area	1) 52.52 sq.m
(6) When the levy or connection is given.	
(7) If the name of the party giving the name / address of the document or the order or order of the Civil Court, the name and address of the reply.	1): Name:[redacted]Age:[redacted] Address:[redacted]
(8) Name and address of the respondent, if there is a decree or order of the parties,	1): Name:[redacted]Age:[redacted] Address:[redacted] 2): Name:[redacted]Age:[redacted] Address:[redacted]
(9) Date of the date of the document	16/01/2017
(10) Date of registration of the document	16/01/2017
(11) Serial numbers, Volumes and Pages	249/2017
(12) Stamp duty as per market price	325000
(13) Registration Fee as per marketable	30000
(14) Remarks	

Details taken for the assessment: -

Selected article on stamp duty: -: (i) In the limits of any Municipal Corporation or any Cantonment area to annexed it.

Figure B10
Translated Transaction Document: An Example

This figure presents the translated version of the original document in Figure B9 using Google Translation services.

The cross-sectional distribution is also large. At the lowest end of the distribution are sub-zones with ₹7330 as the circle rate per square meter of property area (in 2014), to ₹653,240 per square meter of property area in 2018. Figure B11 presents the geo-spatial variation at the sub-zone level in circle rates at the beginning of our sample (Panel A) and at the end of the IGR sample in 2018 (Panel B). The circle rates have been rescaled to the mean sub-zone, with the darker red indicating sub-zones with high circle rates and sub-zones in lighter shades of yellow indicating those with the lowest circle rates in Greater Mumbai.

Table B3
Circle Rate - Summary Statistics

Year	Sub-zones # (1)	% with Change (2)	% Change (3)	Cross-sectional Distribution (×1000₹) Mean (4)	1% (5)	25% (6)	50% (7)	75% (8)	99% (9)
2014	727	-	-	139.59	7.33	81.35	109.00	161.85	580.89
2015	747	100.00	14.44	160.21	14.72	94.60	126.20	189.55	619.25
2016	734	97.49	10.54	172.75	30.04	103.92	134.45	201.05	652.03
2017	734	68.46	6.98	178.40	11.62	109.78	145.15	209.88	653.24
2018	734	0.00	0.00	178.40	11.62	109.78	145.15	209.88	653.24

This table reports the summary statistics on variation in circle rates across sub-zones in Greater Mumbai. Column 1 reports the total number of sub-zones in each year of our sample. Column 2 reports the percentage of sub-zones that witnessed a change in circle rates compared to the previous year. Column 3 presents the average percentage change in circle rates relative to the previous year. Columns (4–9) present the cross-sectional distribution of circle rates in 1000s of rupees per square meter of property area.

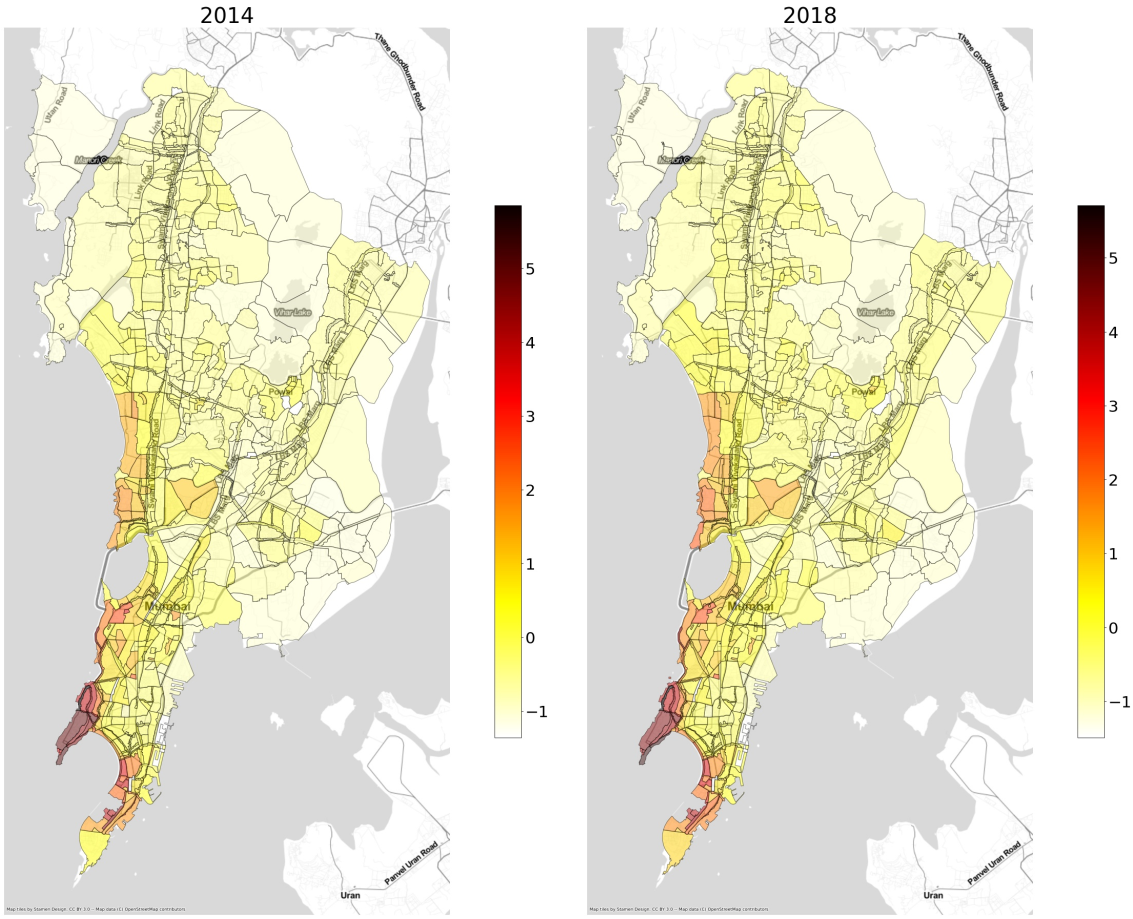

Figure B11
Sub-Zones

Panel (a) presents a heatmap of the sub-zones in 2014, and Panel (b) in 2018. The circle rates are rescaled to the mean sub-zone with darker shades representing the sub-zones with the largest circle rates (in Southern Mumbai) and subzones in white the lowest (northern periphery of Greater Mumbai).

C Measurement Error: Additional Analyses

C.1 Dealing with Classical Measurement Error: Aggregation

In this Section, we describe how aggregating tax base values ($\max[r, c]$) and our proxy market value p help address classical measurement error. We do this by simulating a simple model of under-reporting as follows. Let $f^j(m)$ be the probability density function of market values for all households of type j, where $j \in \{$under-reporter,truthful-reporter$\}$. Observing household $i's$ triple (m, r, c) reveals its (unobservable) type perfectly. If $r = c$ and $m > r$, then $j =$ under-reporter. If $r = m$, then $j =$ truthful reporter. θ (the share of under-reporters) is then simply the fraction of all households with $r = c$ and $m > r$. The *aggregate* amount of under-reporting is the difference between aggregate m and aggregate r for households with $r = c$, because aggregate under-reporting for truthful reporters equals zero, and all under-reporters bunch at $r = c$. At the bunch point c, under-reporters are well identified.

Figure C12a simulates the distributions of r and m around c assuming that $\theta = 40\%$, $(m - c) \sim \mathcal{N}(\mu = 1, \sigma^2 = 10)$, and $c \sim \mathcal{N}(\mu = 10, \sigma^2 = 1)$. The x-axis is $\frac{r-c}{c}$ for the r distribution, and $\frac{m-c}{c}$ for the m distribution. The figure reveals substantial bunching of r around c, with 40% of households with $m \geq c$ choosing to report c. The underlying θ parameter can be backed out by inspecting how the (bunched) distribution of r around c differs from the (smoother) distribution of m around c.

As mentioned, the extent to which the government-assessed values c track market values m is a key confound. To evaluate this issue, we need an independent measure of m which is neither r nor c. For now, Figure C12a is drawn under the assumption that we have access to a measure p which is a perfect estimate of m.

The first approach to distinguish the case with measurement error from genuine under-reporting is to compare *aggregated* reported and market values by $\frac{r-c}{c}$ bins. Figures C12b and C12d plot simulated aggregated reported and market values for the two cases considered above (in Figures C12a and C12c respectively). In these figures, we sum all measures of value (p, r, c, m) for the simulated transactions within $\frac{r-c}{c}$ bins of 0.02 width. In both figures, the x-axis indicates the total reported value of transactions; for example, in Figure C12b at zero, the blue bar is the sum of all r for transactions with $-0.02 \leq \frac{r-c}{c} < 0.02$, and the green bar is the total market value of transactions in each $\frac{r-c}{c}$ bin. In Figure C12d, we separately plot red bars, which show the sum of all p within each bin.

In Figure C12b, with simulated high under-reporting, but no measurement error in p, we can clearly see under-reporting in the zero bin, where aggregated r is substantially lower than aggregated m (since there is no measurement error assumed, the sum of all m also equals aggregate p). In contrast, in the simulation in which there is no under-reporting, but substantial measurement error in p, Figure C12d shows that aggregated m and aggregated noisy p are (approximately) the same within each bin. The insight here is that aggregation within bins smooths out symmetric measurement error in p, and allows us to observe differences between r and m. This allows us to distinguish between the two cases, which are identical in the simulated count/bunching distributions plotted in Figures C12a and C12c. Put differently, truthful reporting with measurement error will exhibit no mass difference between aggregated p and r, and under-reporting will result in a mass difference between aggregated p and r in the central bin.[43] Of course, the data is likely to feature a mix of under-reporting and measurement error in the proxy p for m, and Figure C12f shows that in the presence of both under-reporting and measurement error, aggregation within bins smooths out measurement error and still facilitates a direct comparison of the average differences between r and m.

Figure 2b applies the aggregation procedure to the actual data, plotting aggregated reported and Propequity values in the data by $\frac{r-c}{c}$ bins. The green triangles indicate the total amount of Propequity estimated value (p) transacted amongst transactions within a 0.02 width $\frac{r-c}{c}$ bin, and the blue dots indicate the total value reported (r) for the same set of transactions within the $\frac{r-c}{c}$ bin. The figure is consistent with Figure C12f, i.e., we confirm that there is under-reporting using the aggregation approach, and the figure reveals that the largest amount of unreported value (the differential between aggregated r and p) comes from bunching transactions that report r exactly equal to c. It is also worth noting that in Figure 2b, the overall overlap of the p and

[43] Wider bins can, to a point, help us to smooth measurement error even further and better identify under-reporting at the expense of moving away from the sharp point at $r = c$.

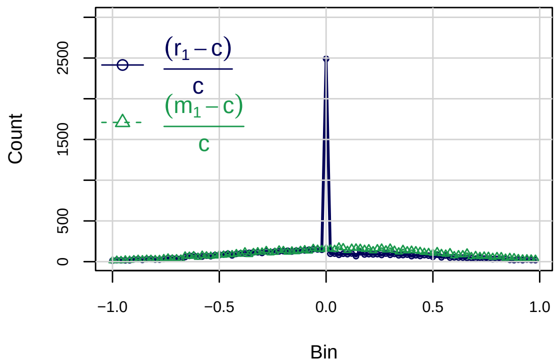

(a) High under-reporting, without measurement error

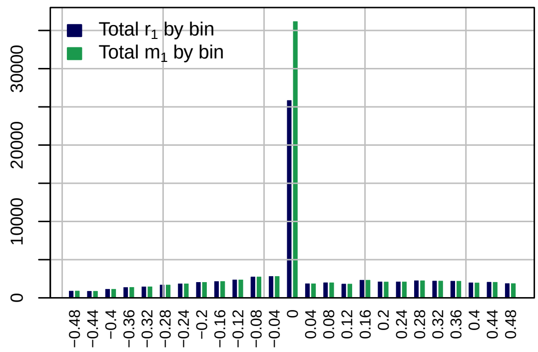

(b) High under-reporting, without measurement error

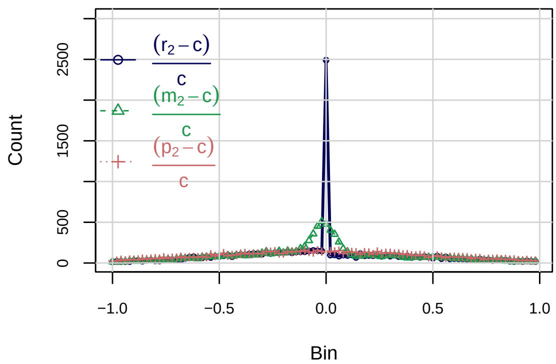

(c) No under-reporting, with measurement error

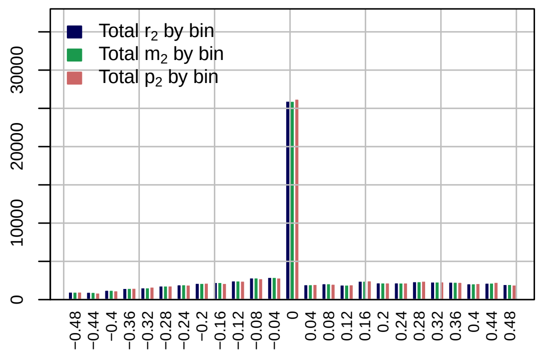

(d) No under-reporting, with measurement error

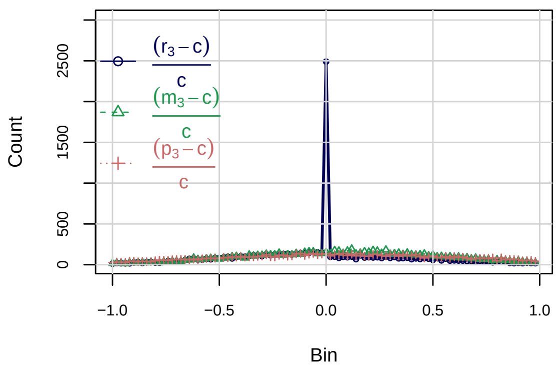

(e) High under-reporting, with measurement error

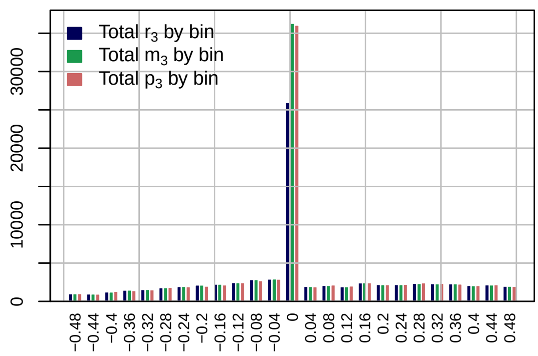

(f) High under-reporting, with measurement error

Figure C12
Simulation Results

r_1 and m_1 are reported and transaction values in the high under-reporting case. r_2, m_2 and p_2 are reported, true transaction, and noisily measured transaction price variables for the no under-reporting case. r_3, m_3 and p_3 are reported, true transaction, and noisily measured transaction price variables for the high under-reporting case with measurement error. c is the guidance value. Measurement error refers to noise in our estimates of market prices relative to the true unobserved market price.

r distributions is tighter than in the plot of transaction counts in Figure 2a, with the sharpest deviation evident between the distributions at exactly c. This pattern is more closely consistent with the "two-type" simplified model presented in the beginning of this section, which predicts an atom of mass in the distribution of r at c and a smooth distribution that is closer to p otherwise.[44]

C.1.1 Measurement Error Correction for $r > m$

In this Section we discuss our measurement error correction approach to the empirical result that the fraction of over-payments is increasing as we observe transactions with larger and larger values of $\frac{r-c}{c}$. The intuition for our measurement error correction is as follows. As we observe transactions with larger $\frac{r-c}{c}$ values, we are looking at transactions that also (likely) have larger m values relative to c and p. Some of that error relative to c and p comes from the error between p and c, but some also comes from the error between p and m; this gives a mechanical reason why our p measure will specifically tend to underestimate m when we focus on cases where m is large relative to c.

Formally, this measurement error issue is a type of "shrinkage" problem where our best predictor does mechanically worse as we focus on more extreme outcomes. Suppose our proxy price p is a mean-zero best guess at the true market value m, using some information set X_p, so that $p = E[m|X_p]$, or $m = p + \epsilon_p$ where $E[\epsilon_p|X_p] = 0$. Suppose further that the guidance value c is also a noisy estimate of m, but uses less information than p, so $c = E[m|X_c]$ where X_c is a subset of information X_p. Then we can write $c = E[p|X_c]$ or $p = c + \epsilon_c$, where $E[\epsilon_c|X_c] = 0$. Combining, we have $m = p + \epsilon_p = c + \epsilon_c + \epsilon_p$. Note that $m - p = \epsilon_p$, so $E[m - p|m - c] = E[\epsilon_p|\epsilon_c + \epsilon_p]$. If ϵ_c and ϵ_p are both normal, then $E[\epsilon_p|\epsilon_c + \epsilon_p] = (\sigma_p^2)/(\sigma_p^2 + \sigma_c^2)(\epsilon_c + \epsilon_p)$, i.e. the ϵ_p term we observe will be a "shrunk" version of the true gap between c and m.

An appealing feature of this measurement error framework is that although there is no obvious reason why the share of households reporting $r > m$ (i.e. over-reporting) should be increasing with $m - c$, the framework directly predicts that the fraction of

[44] For bins with $\frac{r-c}{c} < 0$ (i.e. transactions with reported values less than guidance values) the blue circles in Figure 2b aggregate c. This is because if $r < c$, the tax base is effectively c, i.e., the guidance value, since the government assesses taxes at c pending a successful appeal.

transactions with $r > p$ will increase as $r - c$ increases (a pattern we observe in the data). The framework suggests the following measurement error adjustment. Let $\hat{\epsilon_p}$ be an adjustment we can make to p to inflate $p - c$ back up to an unbiased estimate of $m - c$ conditional on $r(c)$. Our measurement error correction is to adjust $(p - c)/c = \epsilon_c/c$ to $(\hat{\epsilon_c}/c) = \alpha\frac{\epsilon_c}{c}$. We choose α so that the share of transactions with $r(m) > p$ is stable as we look at higher $r - c$ reports. Intuitively, this removes the tendency of p to be too low because we are looking at "large" error transactions when we look in the $m - c$ space.

Figure C13c shows our measurement error corrected inaccuracy, over-payment and under-payment ratios. The correction leads our over-payment ratio to be flat throughout the whole reporting distribution, which, based on our measurement error theory, suggests we have removed the measurement error problem with $\alpha = 1.9$. After our measurement error correction, we continue to find the pattern that there is very little over-payment for transactions where $\frac{r-c}{c} < 0$ - again suggesting guidance values are set relatively low compared to market values.

Figure C13d directly shows the fraction of estimated market value under-reported by $\frac{r-c}{c}$ bin. We present both under-reporting rates based on our raw p measures and our error-corrected $\hat{m}$ measures. Under-reporting is calculated as the $p - \max[r, c]/p$ for the blue curve and $\hat{m} - \max[r, c]/\hat{m}$. Both series show that those reporting $\frac{r-c}{c} < 0$ tend to have high *under-reporting* rates - again consistent with the government setting guidance values quite low to avoid over-payments. As expected, both series show a peak in under-reporting near $r = c$. The error-correction leads us to estimate approximately a stable 30% under-reporting rate for transactions where $\frac{r-c}{c}$ increases. While our main goal is not to estimate the aggregate under-reporting rate, we note that it ranges from 12% under the assumption of no measurement error to 31.7% under the measurement error correction with $\alpha = 1.9$.

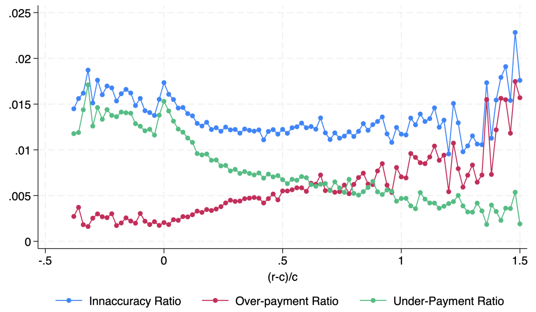

(a) Tax Collected Inaccuracy Ratios without Measurement Error Correction

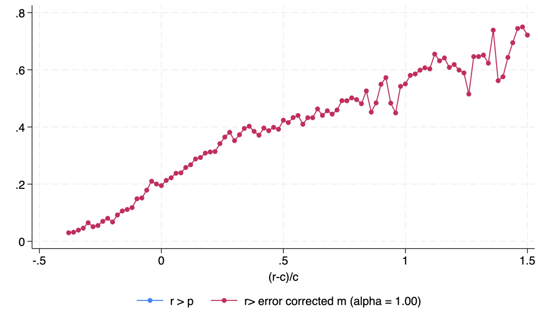

(b) Fraction of Transactions Reported Value r $>$ Propequity Value p

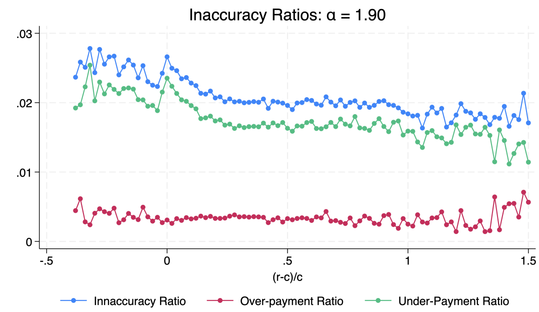

(c) Tax Collected Inaccuracy Ratios with Measurement Error Correction ($\alpha = 1.9$)

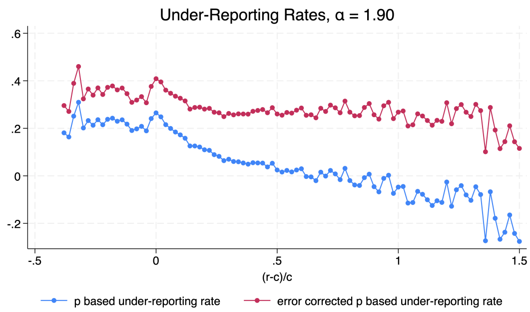

(d) Under-Reporting Rate With and Without Measurement Error Correction

Figure C13

Tax Inaccuracy Ratios Before and After Measurement Error Correction

Panel (a) shows the aggregate inaccuracy ratio within $\frac{r-c}{c}$ bins. A bin's aggregate inaccuracy ratio is defined as the sum of the absolute values of the differences between the tax paid and the tax rate times our proxy market value p, divided by the total p in a bin. The Over-Payment Ratio is the same, but only includes over-payments (i.e. cases where the tax paid is larger than the tax rate times p). The under-payment ratio is the same, but only including under-payments. Panel (b) shows the fraction of transactions where $r > p$. Panel (c) shows the ratios from panel (a) after the measurement error correction described in the Appendix. Panel (d) shows the under-reporting rates using the raw Propequity p data and the error-corrected p data.

Figure C14 reproduces our main empirical results in Figure 2, but substitutes the error-corrected market value proxy for m. The main difference is a rightward-shifted market value distribution (which results from $\alpha = 1.9$ error correction parameter). Under this error correction, the model-implied reported value distributions (in terms of counts in Figure C14a and values in Figure C14b) continue to match the true reported value distributions. The model is also able to match the higher under-reporting rate of approximately 30% at bins .2 and higher in Figure C14c. Re-estimating our ψ parameter using the same methodology in Section 5, we estimate $\psi =$ ₹2.3 compared to the nonerror corrected $\psi = ₹2.12$. The government's preference for tax accuracy is ultimately similar with or without this correction.

C.2 Exact Matching of Propequity Projects to Transactions

A different way to assess the importance of measurement error in p is to focus on a subset of the data that has plausibly less measurement error. Appendix Figure C15 presents bunching results based on a restricted dataset of the 60% of transactions where there is an exact match between building names in the Propequity data and the Registrar transactions data. To the extent that measurement error arises from imperfect matching across buildings (e.g., transactions in older buildings being matched to new luxury building launches in Propequity), we expect this sample to be less affected by measurement error. The bunching patterns appear slightly sharper (bins just to the right of zero do not show as much increase in mass as the bins just to the right in the full sample version. The shape of estimated levels of under-reporting by bin is essentially the same (Appendix Figure C15c versus Figure 2c), but the exact match sample levels are lower in some bins.[45] Overall, the aggregate under-reporting rate in the exact match sample is 9.48% (95% bootstrapped C.I. = [9.37%,9.48%]), which is similar to the full sample estimate of 10.94% (95% bootstrapped C.I. = [10.8%, 11.31%]).

[45] For example, in the exact sample match the estimated under-reporting rate when $\frac{r-c}{c} = 0$ is approximately 0.2, whereas it is 0.25 in the full sample.

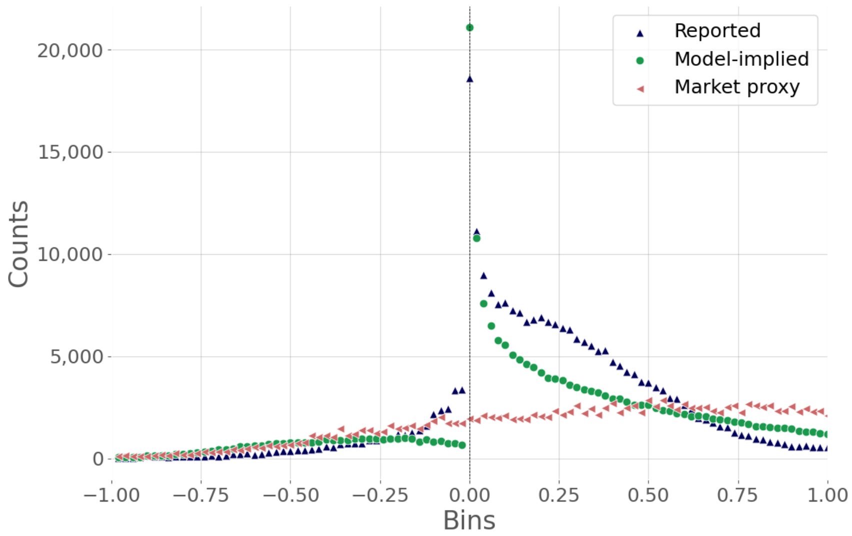

(a) Count Distribution

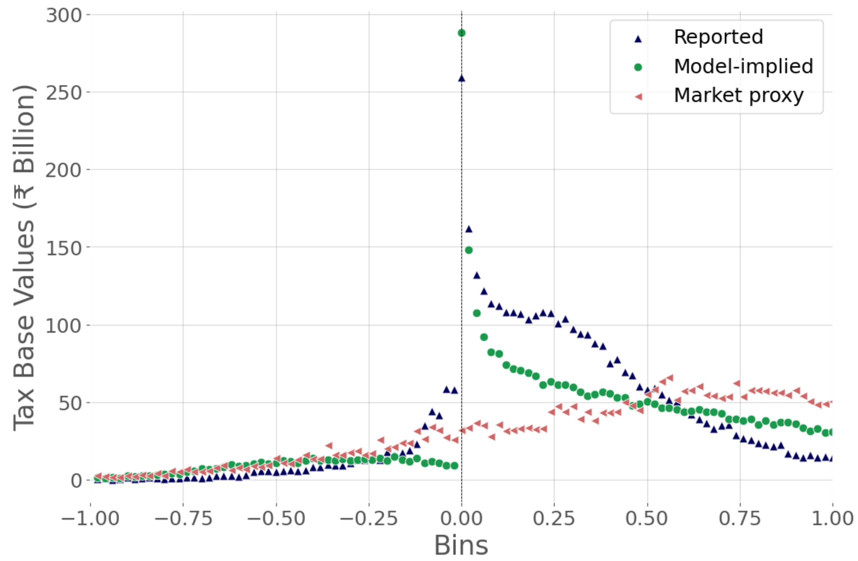

(b) Value Distribution

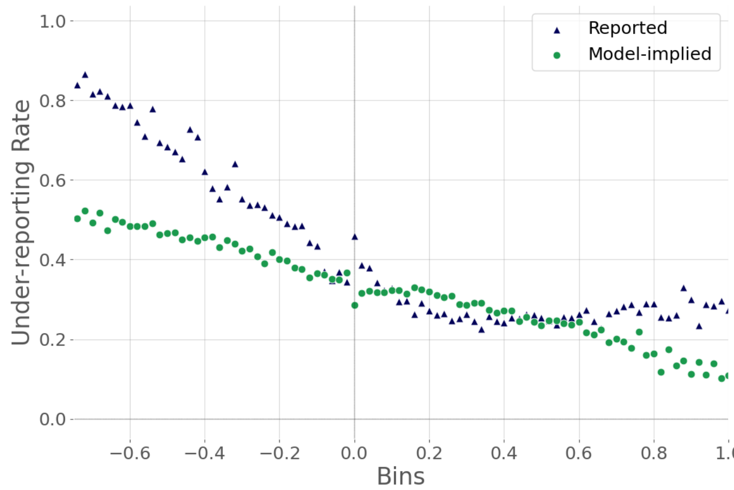

(c) Under-reporting Rate

Figure C14

Raw Data and Structural Model Fits with Measurement Error Correction

This figure shows the raw reported value data, measurement error-corrected market proxy values, and structural model fits under the measurement error method developed in the Appendix. Panels (a) and (b) present the model fit to the targeted moments, the distribution of transaction counts in 2% bins and the measurement error-corrected market proxy value distribution. Panel (c) presents the model fit to the untargeted under-reporting rate, where under-reporting rates are calculated using the measurement error-corrected market values.

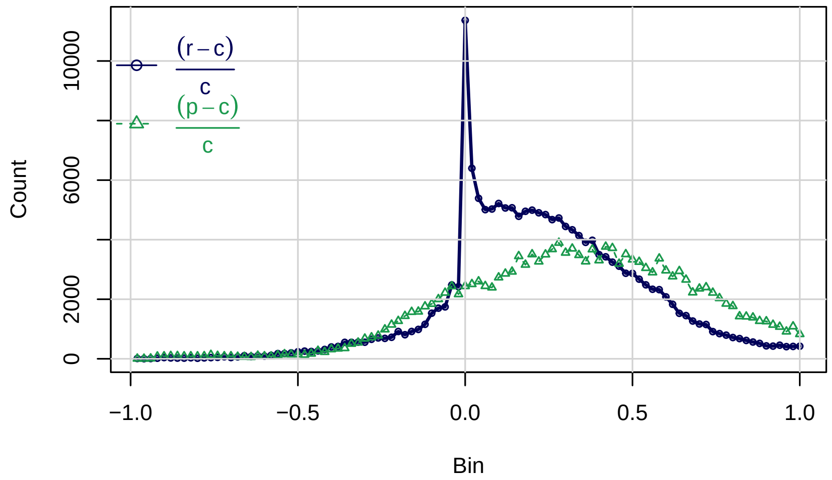

(a) Exact Match r and p Distributions

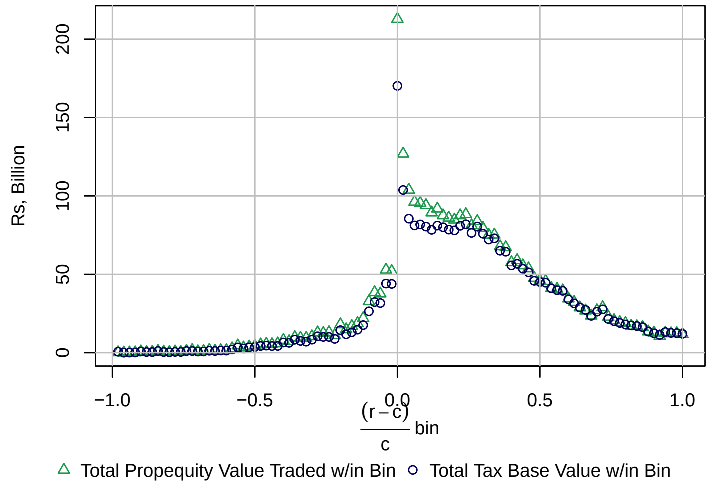

(b) Exact Match Aggregate Tax Base Distribution

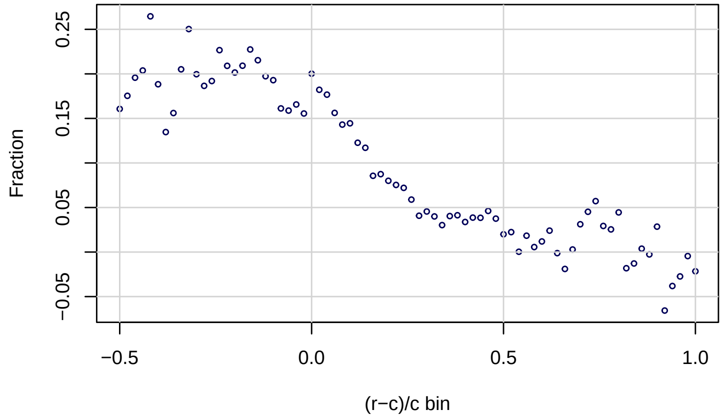

Figure C15

Exact Match Results on Bunching, Aggregate Values, and Under-Reporting Rates

r = reported value, c = guidance value, p = our estimate of market values. This figure shows r and p distributions around c using the 60% of transactions in our sample where we have an exact project match for our estimate of the market price (the Propequity values). Panel (a) shows the raw data distribution of r around c, and p around c. Panel (b) shows the aggregated reported value within the 2% $\frac{r-c}{c}$ bins. Panel (c) presents the estimated under-reporting rate within the 2% bins.

C.3 Stale Guidance Values

For measurement error to explain our results rather than under-reporting, c must perfectly track m across both time and space to explain visible bunching of r at c. However, c is set in a geographically coarse manner and infrequently updated. Over the sample period, real estate prices have grown substantially, with growth rates that vary considerably in different Mumbai sub-regions, making it implausible that c perfectly tracks m over time, given the government's process for setting c. Moreover, c is set at a relatively broad geographical level. Given the considerable spatial variation of property values, it is also implausible that c perfectly tracks m spatially. In the data, c per square meter values are very close or the same for all properties within each sub-zone (the average (median) subzone in the data is 686,818 (264,136) sqm).[46] This means that a single guidance value for a large region is unlikely to be an accurate reflection of the full distribution of the true value of the assets in that region at any given point in time.

Despite these issues, it might still be the case that c values are set carefully to match the first spatial moment of m. However, if this is true, the mass of transactions happening at prices above (infrequently updated) c will rise or fall over time as house prices grow or shrink on average, given regional and aggregate price variation. This generates a concrete prediction. If reporting is truthful, with such time-variation we would expect to find the greatest bunching of r at c immediately *after* c values are updated (i.e., when c is closest to m), and a gradual decline in bunching as m drifts away from c before c is updated again. If counterparties anchor m at government-determined c, a similar prediction obtains, as the accuracy and relevance of c might be expected to be highest immediately after it is updated. Moreover, infrequent updates in the presence of anchoring can create incentives for sellers to wait for c values to increase, as it could allow them to negotiate for substantially higher prices (in the three years we observe guidance value changes the average increase were 14.4 %, 10.5 %, and 6.98 %, see Table B3).[47]

[46] The guidance values do incorporate some adjustments for whether a building is categorized as luxury, the floor the apartment is located in, and whether a parking space is included; but these are all categorical adjustments that are essentially swamped by the price variation within locations across buildings.

[47] Even if sellers believe demand will be lower after c increases, they should be able to obtain some of the surplus generated by transacting at lower c values in the present by waiting and transacting at

In contrast, if buyers and sellers under-report to evade taxes, and with the dates of c changes publicly announced in advance, we would expect that bunching would be greatest immediately *before* any increases in the government-assessed value. If c predictably increases, there is a predictable jump in the tax burden incurred by under-reporting the transaction value at $r = c$ immediately after the rise in c as opposed to immediately before the rise in c, delivering a strong incentive to under-report prior to the change in c.

In the data, Figure C16a documents behavior that is consistent with the under-reporting explanation, and inconsistent with either the measurement error or anchoring explanations. The figure shows that bunching mass at $r = c$ spikes immediately *prior* to scheduled guidance value changes but there are no corresponding spikes in the third-party estimated proxy p.[48] While this is clearly evident in the plot, it is difficult to draw strong conclusions relative to the month-by-month variation overall from pure visual inspection. Table A1 therefore estimates a regression model in which we explain the under-reporting rate using a time-trend and month of year fixed effects, and check whether the under-reporting rate varies in months prior to assessed value changes. The table shows that under-reporting rates are indeed approximately 6 % higher (statistically significant at the 5% level) in months prior to changes in assessed values c. This is a large increase in under-reporting relative to the sample average under-reporting rate of 6%.[49]

Finally, more support for the under-reporting explanation is provided by Figure 3b, which uses the observed agreement date and the registration date for each transaction in the data to check for backdating behavior. The figure shows that there is a clear pattern of backdating agreement dates just prior to guidance value changes to take

higher future c values tomorrow. Such arguments depend on discount rates and demand and supply elasticities, and there are possibly constellations of parameter values that can deliver greater bunching prior to c value changes under truthful reporting. Ultimately, given further evidence described below, Occam's razor suggests that such arguments are potentially less plausible than under-reporting being higher amongst buyers who backdate transaction times to take advantage of infrequent updates.

[48] The guidance values were increased in 2014, 2015 and 2016. The Maharashtra government chose not to increase the guidance values in 2017, 2018, 2019 and 2020, and decreased the rate as a result of the pandemic in 2021.

[49] This sample average is the simple average under-reporting rates across transactions; this differs from our aggregate under-estimating estimate of 10.94% because it does not weight transactions by their size.

advantage of lower guidance values.

C.4 High-Frequency (Daily) Reporting Behavior

Figure C16 studies *daily* reporting behavior around scheduled guidance value changes.[50] Figure C16 shows a large spike in registered transactions on the day directly before the scheduled guidance value change and Figure C16b shows the fraction of transactions registered below the guidance value which also shows an increase in the days prior to the guidance value change. These results suggest that even non-bunching transactors prefer to avoid the new guidance values which are expected to increase their tax burden. This could happen, for example, if the new guidance values are above what buyers planned to report.

[50] The dates of scheduled guidance value changes have moved around over time and are not always on 1 January each year. This reduces concerns that these dates coincide with policy announcements that are routinely made at the turn of the year.

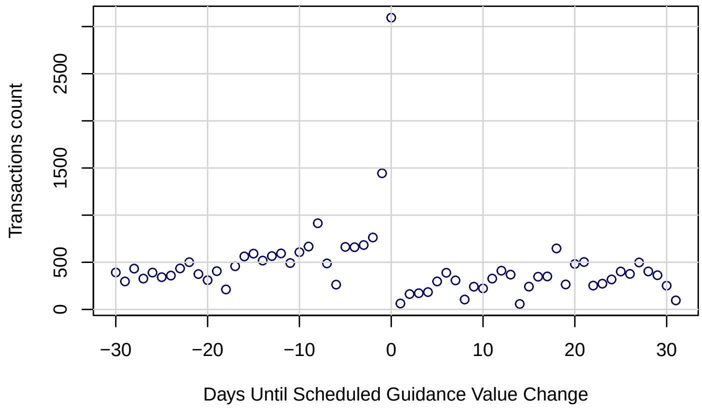

(a) Transaction Counts

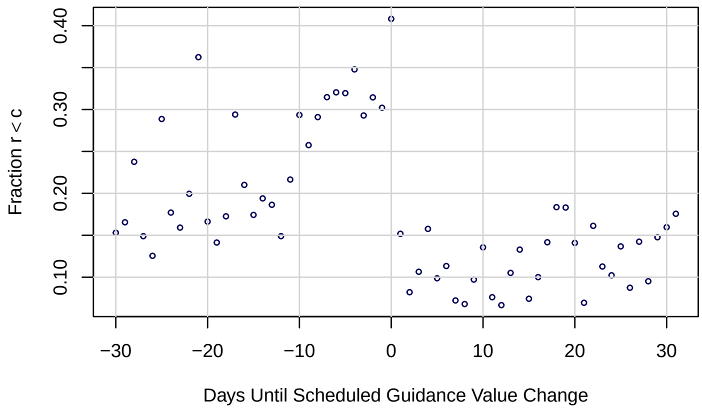

(b) Fraction $r < c$

Figure C16

Reporting Behavior in Days Around Scheduled Guidance Value Increases

Panel (a) shows the fraction of the transactions on the given event-day that "bunchers", i.e. had a reported value within 1% of the guidance assessed value. Panel (b) shows the fraction of transactions that had a reported value less than the guidance assessed value.

D Estimating the Response of Transactions to Changes in Guidance Values

We estimate the extensive-margin response of housing transactions to legislated guidance value (GV) changes. GVs in Mumbai change annually, with variation across subzones and time. Figure D17a summarizes GV changes and transaction patterns. Panel (a) shows average GVs, with large increases in 2013–2015 and smaller ones in 2021–2022. We exploit citywide and subzone-level variation using two designs: a difference-in-differences (DiD) event study and a triple-difference specification.

D.1 Design 1: Difference-in-Differences Event Study

We define treated subzones as those with $>$10% GV increases in a given year; others are controls. We pool monthly subzone observations in a $\pm$9 month window around each GV change. Panels (b)–(c) of Figure D17 show: (i) legislated GVs in treated subzones rise by $\sim$18%, versus $\sim$0% in controls; (ii) transactions' *reported* GVs rise by $\sim$10%. Panel (d) shows no post-policy drop in transaction counts; both groups spike before the change, consistent with timing responses but not sustained declines.

D.2 Design 2: Triple-Difference

The difference-in-difference analysis shown so far does not exploit the full distribution of guidance value variation, and also does not exploit the fact that we expect transactions whose reported value is closer to the guidance value to be affected more by changes to guidance values. We now present results from a triple-difference design that exploits this variation.

Let c and c' be the pre- and post-change guidance values (GV) per square foot in a subzone with proportional change g. If the extensive-margin elasticity is high, we would expect transactions reporting below c' to fall sharply post-change, while those already reporting above c' remain stable, as they face no new binding constraint. To pool subzones with different c and c', we define a standardized measure $\bar{r}$. For pre-change transactions, $\bar{r} = \frac{r-c}{\hat{c}'-c}$, where $\hat{c}' = (1+g)c$; for post-change transactions, $\bar{r} = \frac{r-\hat{c}}{c'-\hat{c}}$, where $\hat{c} = c'/(1+g)$. By construction, $\bar{r} = 1$ corresponds to reporting exactly at the post-change GV; $\bar{r} < 1$ (treated bins) captures transactions likely affected by the change, while $\bar{r} \geq 1$ (control bins) are unlikely to be affected.

Figure D18a plots histograms of $\bar{r}$ before (blue) and after (red) the policy change: pre-change bunching at $\bar{r} = 0$ reflects transactions reporting exactly the old GV, while post-change bunching at $\bar{r} = 1$ reflects reporting exactly the new GV. Under a high elasticity, we would expect little red mass below $\bar{r} = 1$ and the red distribution to mirror the blue above 1; under a zero elasticity, most blue mass below $\bar{r} = 1$ would shift to $\bar{r} = 1$. The data more closely match the latter, suggesting a low elasticity. However, this figure does not exploit across-subzone variation in the magnitude of GV changes, which we turn to next in the following triple-difference regression:

$$y_{ict} = \beta_0 + \beta_1 P_t + \beta_2 T_{ic} + \beta_3 \Delta g_i + \beta_4 P_t T_{ic} + \beta_5 P_t \Delta g_i + \beta_6 T_{ic} \Delta g_i + \beta_7 P_t T_{ic} \Delta g_i + \epsilon_{ict}$$

where y_{ict} is the transaction count in subzone i, treatment/control portion of the reported value distribution $\bar{r}$, event-month t; P_t is a post-change dummy; T_{ic} is a treatment dummy; and Δg_i is the subzone's GV change. β_7 measures the differential post-change response in treated bins in subzones with larger GV increases.

Figure D18 shows the slope of transaction counts vs. GV changes by treatment/control and pre/post, for treatment definition cut-offs of 1.1× to 1.9× GV (we alter treatment definition cut-offs to allow for possible impacts of guidance value changes above c'). Slopes are similar across all groups, suggesting small or zero elasticities. Figure D19(a) plots β_7 across different potential treatment definition cut-offs - the coefficient asymptotes to close to zero suggesting little differential change in quantity of transactions. Panel (b) shows corresponding elasticities:

$$\epsilon = \frac{\Delta Q/Q}{\Delta GV/GV} \quad \text{where} \quad \Delta Q/Q \text{ is the triple-diff estimate for treated bins post-change.}$$

Even the most negative point is about -0.4, with CIs generally including zero.

Figure D17
Policy Changes, Reported Values, and Transaction Counts

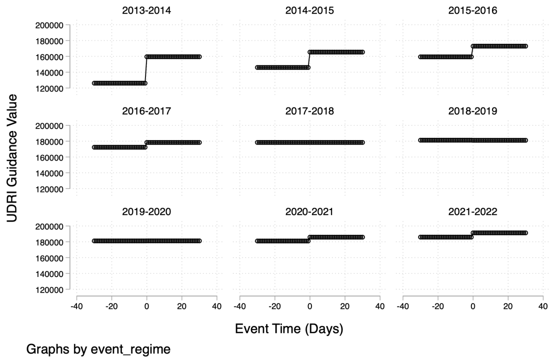

(a) Average GV over Time

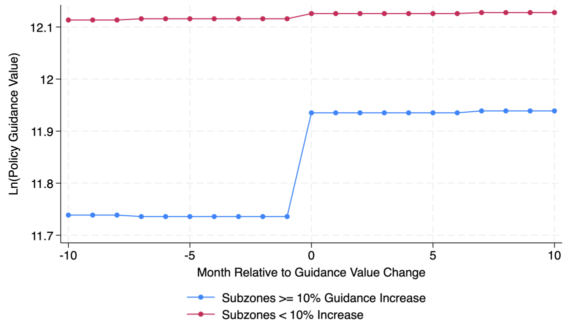

(b) Legislated GV

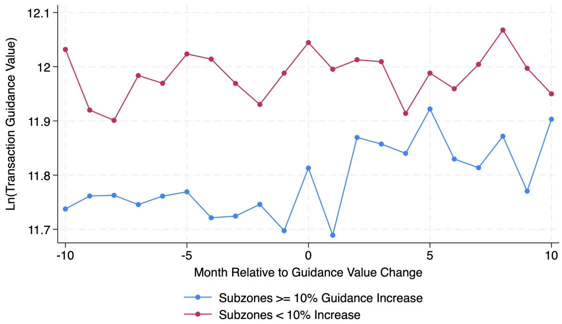

(c) Reported GV

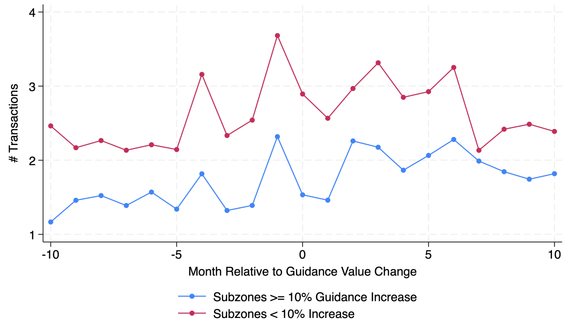

(d) Transaction Counts

(a) Average GV over time. (b) Event-study effects on legislated GV. (c) Effects on transactions' reported GV. (d) Effects on transaction counts; no post-policy divergence is evident.

Figure D18
GV Changes, Treatment/Control Regions, and Transactions by Cut-off

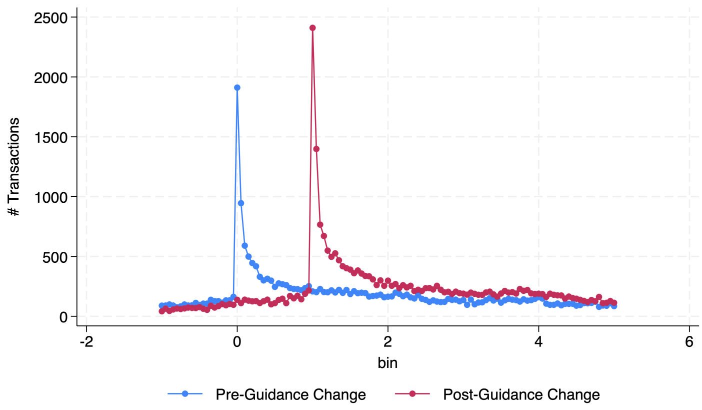

(a) Treatment/Control Definition

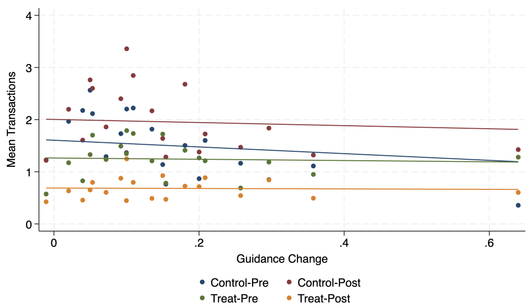

(b) 1.1× GV

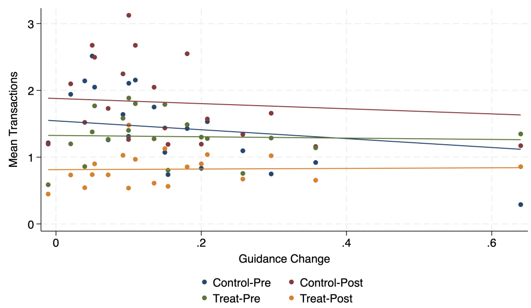

(c) 1.3× GV

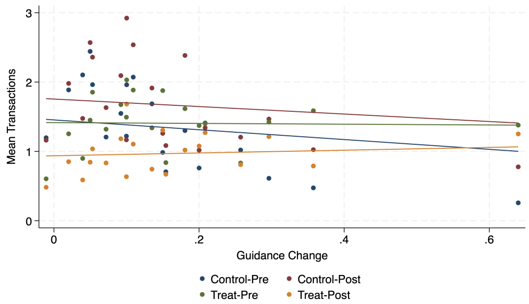

(d) 1.6× GV

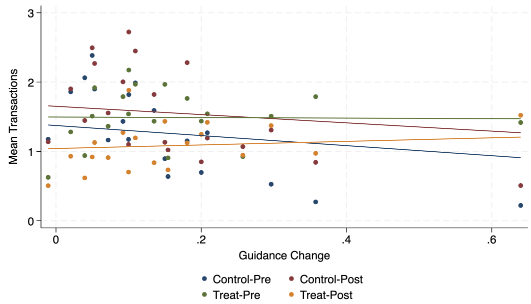

(e) 1.9× GV

Panel (a) shows the definition of treatment/control bins in the standardized reported value distribution for the triple-interaction event-study; Panels (b)–(e) show slopes of transaction counts vs. guidance value (GV) changes before and after the policy change for different cut-offs. Slopes are similar for treated and control bins at all cut-offs.

Figure D19
Triple-Difference Coefficients and Elasticity Estimates

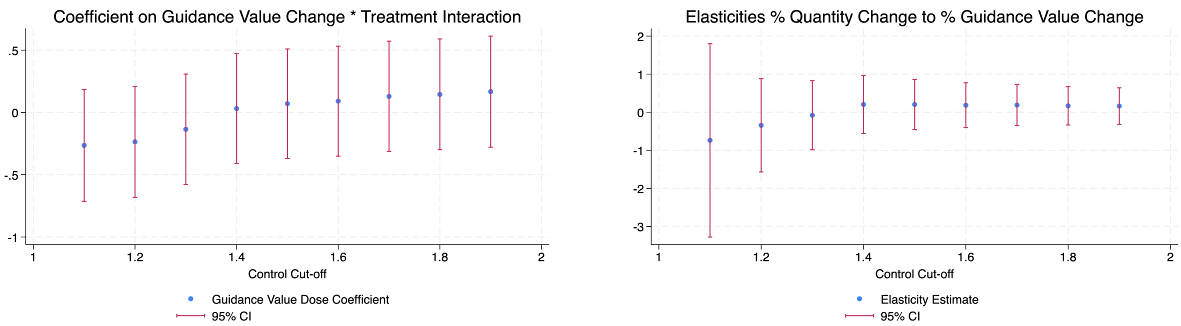

(a) β_7 Coefficients **(b)** Elasticity Estimates

(a) β_7 estimates from triple-difference regressions. (b) Corresponding elasticities; CIs generally include zero.